

NVM Express: SCSI Translation Reference

Please send questions or comments to info@nvmexpress.org

Maintained by the NVM Express Workgroup

Revision 1.5

June 24, 2015

Table of Contents

Introduction	5
1.1 Scope and Theory of Operation	5
1.2 Normative References	6
1.3 Terminology	6
2 Overview of Command Mappings	7
3 Common SCSI Field Translations	9
3.1 ALLOCATION LENGTH	9
3.2 BYTCHK	9
3.3 CONTROL	9
3.4 DPO	9
3.5 GROUP NUMBER	9
3.6 FUA	9
3.7 FUA_NV	9
3.8 LOGICAL BLOCK ADDRESS	9
3.9 PARAMETER LIST LENGTH	10
3.10 PRODUCT IDENTIFICATION	10
3.11 PRODUCT REVISION LEVEL	10
3.12 IMMED	10
3.13 T10 VENDOR IDENTIFICATION	10
3.14 TRANSFER LENGTH	10
3.15 VENDOR SPECIFIC IDENTIFIER	10
3.16 VERIFICATION LENGTH	10
3.17 VRPROTECT	10
3.18 WRPROTECT	11
4 SCSI Primary Command Mappings	12
4.1 INQUIRY (SPC-4)	12
4.2 LOG SENSE (SPC-4)	13
4.3 MODE SELECT(6), MODE SELECT(10) (SPC-4)	14
4.4 MODE SENSE(6), MODE SENSE(10) (SPC-4)	15
4.5 REPORT LUNS (SPC-4)	17
4.6 REQUEST SENSE (SPC-4)	18
4.7 SECURITY PROTOCOL IN (SPC-4)	19

4.8	SECURITY PROTOCOL OUT (SPC-4).....	20
4.9	SEND DIAGNOSTIC (SPC-4).....	20
4.10	START STOP UNIT (SPC-4).....	20
4.11	TEST UNIT READY (SPC-4).....	22
4.12	WRITE BUFFER (SPC-4).....	22
4.13	PERSISTENT RESERVE IN/OUT (SPC-4).....	23
4.13.1	PERSISTENT RESERVE IN	24
4.13.2	PERSISTENT RESERVE OUT	25
5	SCSI Block Command Mappings.....	28
5.1	COMPARE AND WRITE (SBC-3)	28
5.2	FORMAT UNIT (SBC-3).....	29
5.2.1	FORMAT UNIT Parameter list	30
5.2.1	Formatting with Protection Information	30
5.3	READ(6), READ(10), READ(12), READ(16) (SBC-3)	31
5.4	READ CAPACITY(10), READ CAPACITY(16) (SBC-3)	32
5.5	SYNCHRONIZE CACHE(10), SYNCHRONIZE CACHE(16) (SBC-3)	32
5.6	UNMAP (SBC-3).....	33
5.6.1	UNMAP Parameter List.....	33
5.6.2	UNMAP Block Descriptor.....	33
5.7	VERIFY(10), VERIFY(12), VERIFY(16)(SBC-3).....	33
5.8	WRITE(6), WRITE(10), WRITE(12), WRITE(16) (SBC-3).....	35
5.9	WRITE LONG(10), WRITE LONG(16) (SBC-3)	35
6	SCSI Data	37
6.1	INQUIRY Data	37
6.1.1	Standard INQUIRY Data Page	37
6.1.2	Supported VPD Pages Data Page.....	38
6.1.3	Unit Serial Number VPD Page	39
6.1.3.1	PRODUCT SERIAL NUMBER and PAGE LENGTH fields.....	39
6.1.3.1.1	NGUID-based Translation	39
6.1.3.1.2	EUI64-based Translation	39
6.1.3.1.3	V1.0-compliant Devices.....	40
6.1.4	Device Identification Data Page.....	40
6.1.4.1	NAA IEEE Registered Extended designator format	41

6.1.4.2	NAA Locally Assigned Designator Format	42
6.1.4.3	T10 Vendor ID Based Descriptor	43
6.1.4.4	SCSI Name String Designator	44
6.1.4.4.1	SCSI NAME STRING field.....	44
6.1.4.4.1.1	NGUID-based Translation	44
6.1.4.4.1.2	EUI64-based Translation.....	45
6.1.4.4.1.3	V1.0-compliant Devices	45
6.1.4.5	EUI-64 based Designator.....	45
6.1.4.5.1	EUI-64 Field	46
6.1.5	Extended INQUIRY Data VPD Page	46
6.1.6	Block Limits VPD Page	51
6.1.7	Block Device Characteristics VPD Page	54
6.1.8	Logical Block Provisioning VPD Page.....	54
6.2	Log Pages.....	57
6.2.1	Supported Log Pages	57
6.2.2	Temperature Log Page	57
6.2.3	Solid State Media Log Page.....	58
6.2.4	Information Exceptions Log Page.....	59
6.3	Mode Parameters	60
6.3.1	Mode Parameter Headers	60
6.3.2	Mode Parameter Block Descriptor.....	61
6.3.3	Mode Page Formats	61
6.3.3.1	Read-Write Error Recovery Mode Page.....	62
6.3.3.2	Caching Mode Page.....	64
6.3.3.3	Control Mode Page	65
6.3.3.4	Power Condition Control Mode Page	66
6.3.3.5	Return all pages Mode Page	67
6.3.3.6	Return all pages and subpages Mode Page	67
6.3.3.7	Informational Exceptions Control Mode Page.....	68
6.4	READ CAPACITY Parameter Data	69
6.5	Sense Data.....	70
6.5.1	Fixed Format Sense Data	70

6.5.2	Descriptor Format Sense Data	70
6.6	REPORT LUNS Parameter Data	71
6.7	PERSISTENT RESERVE IN/OUT Parameter Data	71
6.7.1	PERSISTENT RESERVE IN Parameter Data for READ KEYS	71
6.7.2	PERSISTENT RESERVE IN Parameter Data for READ RESERVATIONS	72
6.7.3	PERSISTENT RESERVE IN Parameter Data for READ CAPABILITIES	72
6.7.4	PERSISTENT RESERVE IN Parameter Data for READ FULL STATUS	73
6.7.5	PERSISTENT RESERVE OUT Parameter List	74
7	Status Mappings.....	78
7.1	Generic Command Status.....	78
7.2	Command Specific Status	79
7.3	Media Errors	79
8	Task Management Functions	81

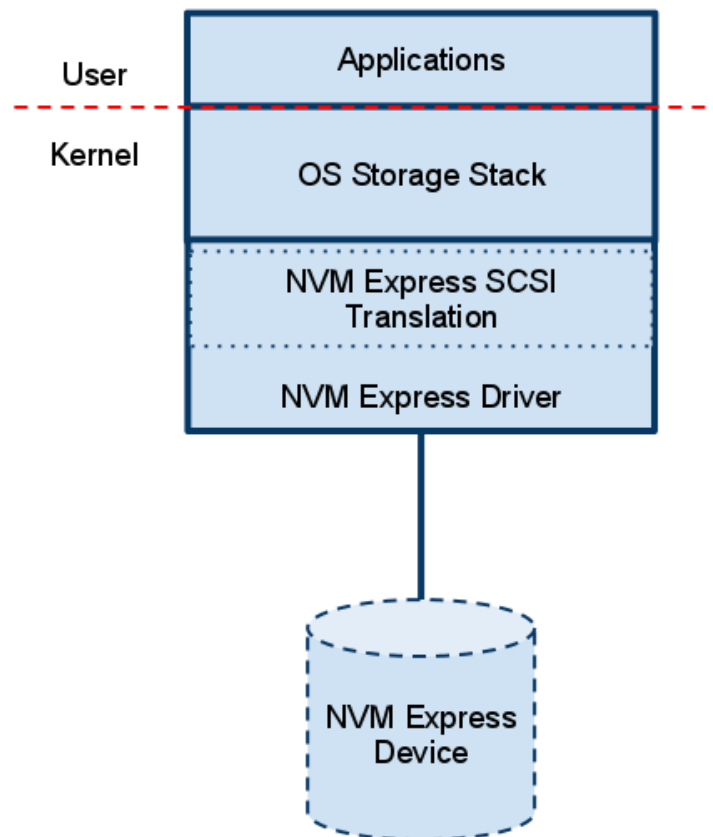
Introduction

1.1 Scope and Theory of Operation

This document defines an NVM Express translation reference for SCSI. NVM Express defines a host controller interface and command set for use with non-volatile memory (NVM) devices, such as PCI Express solid state drives (SSDs). There are existing applications and software infrastructure built upon the SCSI architectural model defined by T10. The purpose of this document is to define a mapping between SCSI and NVM Express specifications to enable a seamless transition to NVM Express by preserving existing software infrastructure investments.

Host software implementing the mapping of SCSI commands into NVM Express commands may be logically viewed as a layer within the operating system stack. Figure 1 depicts the SCSI to NVMe Translation Layer (“SNTL”) pictorially. The NVM Express Driver may implement the translation as an upper layer within the driver.

Figure 1: NVM Express SCSI Translation Driver Model



The translation layer exists logically below the operating system storage stack built upon SCSI and logically above the NVM Express driver implementing the NVM Express host interface. Translation requires the mapping of:

- SCSI commands to NVM Express Commands. Refer to Sections 2, 3, 4, 5.
- NVM Express Status to SCSI Status. Refer to Section 7.

1.2 Normative References

The following standards are referenced within this document:

- ISO/IEC 14776-454, SCSI Primary Commands – 4 (SPC-4) [T10/1731-D]
- ISO/IEC 14776-323, SCSI Block Commands – 3 (SBC-3) [T10/1799-D]
- SCSI Architecture Model – 4 (SAM-4) [T10/1683-D]
- SCSI / ATA Translation – 3 (SAT-3)
- NVM Express 1.0b or later

1.3 Terminology

Unspecified: A term designating that this version of this reference does not specify a translation for a SCSI field. A translation for an unspecified field may be specified by future versions of this reference. Translation of fields marked unspecified shall not conflict with other standards in the set of SCSI standards.

SNL: An acronym for SCSI to NVMe Translation Layer, the functional layer defined in this standard that uses an NVMe device to emulate objects in a SCSI logical unit, providing capabilities defined by SCSI standards (e.g., the SCSI Block Commands (SBC-3) and SCSI Primary Commands (SPC-4) standards).

2 Overview of Command Mappings

SCSI defines commands for an assortment of devices, some of which are not pertinent to NVM subsystems. Figure 2 provides a summary of the SCSI commands that this reference document shall provide a translation for. Translation of all other commands not listed in Figure 2 is unspecified.

Figure 2: NVM Express SCSI Command Mappings

SCSI Command	NVM Express Command(s)	Reference
COMPARE AND WRITE	Compare and Write	5.1
FORMAT UNIT	Format NVM	5.2
INQUIRY	Identify	4.1
LOG SENSE	Get Features, Get Log Page	4.2
MODE SELECT(6)	See Note 1	4.3
MODE SELECT(10)	See Note 1	
MODE SENSE(6)	Identify, Get Features	4.4
MODE SENSE(10)		
READ(6)	Read	5.3
READ(10)		
READ(12)		
READ(16)		
READ CAPACITY(10)	Identify	5.4
READ CAPACITY(16)		
REPORT LUNS	Identify	4.5
REQUEST SENSE	See Note 1	4.6
SECURITY PROTOCOL IN	Security Receive	4.7
SECURITY PROTOCOL OUT	Security Send	4.8
SEND DIAGNOSTIC	N/A	4.9
START STOP UNIT	Set Features, Get Features	4.10
SYNCHRONIZE CACHE(10)	Flush	5.5
SYNCHRONIZE CACHE(16)		
TEST UNIT READY	See Note 1	4.11
UNMAP	Dataset Management	5.6
VERIFY(10)	VERIFY	5.7
VERIFY(12)		
VERIFY(16)		
WRITE LONG (10)	Write Uncorrectable	5.9
WRITE LONG (16)		
WRITE(6)	Write	5.8
WRITE(10)		
WRITE(12)		
WRITE(16)		
WRITE BUFFER	Firmware Image Download, Firmware Image Activate	4.12

Notes:

1. MODE SELECT(6), MODE SELECT(10), REQUEST SENSE and TEST UNIT READY commands do not map 1-to-1 to NVM Express commands; a device driver should translate as specified in the referenced sections.

3 Common SCSI Field Translations

This section specifies the translation of common fields shared by multiple SCSI commands.

3.1 ALLOCATION LENGTH

ALLOCATION LENGTH is specified (SPC-4) as indicating number of bytes allocated by the application client in Data-in buffer for parameter data. Support for this field requires ensuring accesses to Data-in buffer are within the range defined by the ALLOCATION LENGTH.

3.2 BYTCHK

Support requires translation to FUA field of the NVM Express command.

3.3 CONTROL

Bit Field	Notes
NACA (bit 2)	If set to 1b, command shall be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code. (Note: section 6.1.1 requires that the NORMACA bit in Standard Inquiry data be set to 0b).
All others	Unspecified.

3.4 DPO

Support unspecified. Disable page out (DPO) specifies retention characteristics which are not supported in NVM Express.

3.5 GROUP NUMBER

Support unspecified.

3.6 FUA

Support requires translation to FUA field of NVM Express command.

3.7 FUA_NV

Support unspecified. FUA_NV requests that data to be read from non volatile cache.

3.8 LOGICAL BLOCK ADDRESS

Support requires translation to Starting LBA field of NVM Express command.

3.9 PARAMETER LIST LENGTH

PARAMETER LIST LENGTH is specified (SPC-4) as indicating number of bytes allocated by the application client in data-out buffer for MODE SELECT parameter data, which shall include header, block descriptor(s) and mode pages. Support for this field requires ensuring accesses to Data-out buffer are within the range defined by the PARAMETER LIST LENGTH.

3.10 PRODUCT IDENTIFICATION

Shall be set to the first 16 bytes of the Model Number (MN) field within the Identify Controller Data Structure.

3.11 PRODUCT REVISION LEVEL

Shall be set to the last 4 ASCII graphic characters in the range of 21h-7Eh (i.e. last 4 non-space characters) of the Firmware Revision (FR) field within the Identify Controller Data Structure.

3.12 IMMED

CODE	Notes
0b	Shall be supported by returning status after command is complete.
1b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.

3.13 T10 VENDOR IDENTIFICATION

Shall be set to "NVMe" followed by 4 spaces: "NVMe ".

3.14 TRANSFER LENGTH

Support requires translation to Number of Logical Blocks (NLB) field of NVM Express command. Note that for SCSI commands with values of TRANSFER LENGTH greater than 2^{16} , multiple NVM Express commands may need to be issued to satisfy the request.

3.15 VENDOR SPECIFIC IDENTIFIER

This field has a size of 36 bits.

3.16 VERIFICATION LENGTH

Support requires translation to Number of Logical Blocks (NLB) field of the NVM Express command.

3.17 VRPROTECT

Support for VRPROTECT requires setting PRACT and PRCHK fields of the NVM Express command:

If BYTCHK set to zero:

VRPROTECT CODE	PRACT	PRCHK
000b, 001b, 101b	1	111b
010b	1	011b
011b	1	000b
100b	1	100b

If BYTCHK set to one:

VRPROTECT CODE	PRACT	PRCHK
000b	1	111b
001b, 010b, 011b, 100b, 101b	1	000b

3.18 WRPROTECT

Support for WRPROTECT requires setting PRACT and PRCHK fields of the NVM Express command:

WRPROTECT CODE	PRACT	PRCHK
000b	1	000b
001b, 101b	0	111b
010b	0	011b
011b	0	000b
100b	0	100b
All other codes shall result in command termination with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.		

4 SCSI Primary Command Mappings

This section defines the translation of SCSI commands applicable to all SCSI devices into NVM Express commands and structures.

4.1 INQUIRY (SPC-4)

A SCSI INQUIRY command shall be translated into an NVM Express Identify command. The INQUIRY command requests information regarding the device. The type of information to return is indicated in the EVPD and PAGE CODE fields of the INQUIRY command. Table 4-1 specifies the translation of INQUIRY fields.

Table 4-1: INQUIRY CDB Field Translation

INQUIRY Field	Notes and References	
OPERATION CODE	Specified as 12h (SBC-3).	
EVPD (Enable vital product data)	CODE	Notes
	0b	Shall be supported by returning Standard INQUIRY Data to application client, refer to 6.1.1.
	1b	Shall be supported by returning Vital Product Data to application client, refer to PAGE CODE field.
PAGE CODE	CODE	Notes
	00h	Shall be supported by returning Supported VPD Pages data page to application client, refer to 6.1.2.
	80h	May be supported by returning Unit Serial Number data page to application client. Refer to 6.1.3.
	83h	Shall be supported by returning Device Identification data page to application client, refer to 6.1.4.
	86h	May be supported by returning Extended INQUIRY data page to application client, refer to 6.1.5.
	B0h	May be supported by returning Block Limits VPD data page to application client, refer to 6.1.6.
	B1h	May be supported by returning Block Device Characteristics VPD Page to application client, refer to 6.1.7.
	B2h	May be supported by returning Logical Block Provisioning VPD Page to application client, refer to 6.1.8.
	All Others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.

ALLOCATION LENGTH	Refer to 3.1.
CONTROL	Refer to 3.3.

4.2 LOG SENSE (SPC-4)

LOG SENSE is used by application clients to retrieve statistical and operational information.

Table 4-2: Translation of LOG SENSE

LOG SENSE Field	Notes and References	
OPERATION CODE	Specified as 4Dh (SPC-4).	
SP	Code	Notes
	0b	Shall be supported by performing LOG SENSE command. Log parameters shall not be saved.
	1b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
PC	Code	Notes
	01b	Cumulative values shall be returned to application client.
	All Others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
PAGE CODE	Code	Notes
	00h	Shall be supported by returning Supported Log Pages log page to application client. Refer to 6.2.1.
	0Dh	Shall be supported by returning Temperature log page to application client. Refer to 6.2.2.
	11h	Shall be optionally supported by returning Solid State Media log page to application client. Refer to 6.2.3.
	2Fh	Shall be supported by returning Informational Exceptions log page to application client. Refer to 6.2.4.
	All Others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
SUBPAGE CODE	Unspecified.	
PARAMETER POINTER	Unspecified.	
ALLOCATION LENGTH	Refer to 3.1.	
CONTROL	Refer to 3.3.	

4.3 MODE SELECT(6), MODE SELECT(10) (SPC-4)

MODE SELECT is used by application clients to specify mode parameters. The MODE SELECT command fields should be translated per Table 4-3. Support for MODE SELECT(6) and MODE SELECT(10) also requires a device driver to translate mode parameter data as specified in 6.3.

Table 4-3: MODE SELECT Field Translations

MODE SELECT Field	Notes and References	
OPERATION CODE	Specified as (SPC-4):	
	SCSI Command	OPERATION CODE
	MODE SELECT(6)	15h
	MODE SELECT(10)	55h
SP	CODE	Notes
	1b	Any saveable mode pages should be saved. If the pages are not saved, then there is no distinction between current and saved pages
	0b	Shall be supported. There shall be no distinction between current and saved pages.
PF	CODE	Notes
	1b	Indicates that the MODE SELECT parameters following the header and block descriptor(s) are structured as pages of related parameters and are defined in the SPC-4 standard.
	0b	May result in command termination with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
PARAMETER LIST LENGTH	Refer to 3.9.	
CONTROL	Refer to 3.3.	

4.4 **MODE SENSE(6), MODE SENSE(10) (SPC-4)**

MODE SENSE requests that mode parameters be returned to the application client.

Table 4-4: **MODE SENSE CDB Field Translations**

MODE SENSE Field	Notes and References	
OPERATION CODE	SPC-4 specifies this field as:	
	SCSI Command	OPERATION CODE
	MODE SENSE(6)	1Ah
	MODE SENSE(10)	5Ah
DBD	CODE	Notes
	0b	Shall be supported by returning zero or more block descriptors in the returned MODE SENSE data.
	1b	Shall be supported by not returning any block descriptors in the returned MODE SENSE data.
LLBAA	If LLBAA is set to 1 then the LONGLBA field in the parameter data may be set to 1, if LLBAA is set to 0 then LONGLBA shall be set to 0.	
PC	CODE	Notes
	00b	Shall be supported by returning current values in mode pages to application client.
	01b	Shall be supported by returning changeable values in modes pages to application client.
	10b	Shall be supported by returning default values in modes pages to application client.
	11b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
PAGE CODE	CODE	Notes
	01h	Shall be supported by returning Read-Write Error Recovery mode page, refer to 6.3.3.1.
	08h	Shall be supported by returning Caching mode page, refer to 6.3.3.2.
	0Ah	Shall be supported by returning Control mode page to application client, refer to 6.3.3.3
	1Ah	Shall be supported by returning Power Condition Control mode page to application client, refer to 6.3.3.4.
	1Ch	Shall be supported by returning Informational Exceptions Control mode page to application client, refer to 6.3.3.7.
	3Fh	Subpage Code 00h shall be supported by returning all supported mode pages as defined in section 6.3.3.5 .
	3Fh	Subpage Code FFh shall be supported by returning all supported mode pages and subpages as defined in section 6.3.3.6.
	All Others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
SUBPAGE CODE	Unspecified.	
ALLOCATION LENGTH	Refer to 3.1.	

CONTROL	Refer to 3.3.
* MODE SENSE (10) only.	

4.5 REPORT LUNS (SPC-4)

REPORT LUNS (SPC-4) returns a list of logical units to the application client. The 4 least significant bytes shall be set to 0h. Support for REPORT LUNS requires the Identify command and returning to the application client REPORT LUNS Parameter Data as specified in 6.6.

Table 4-5: REPORT LUNS CDB Field Notes

REPORT LUNS Field	Notes and References
OPERATION CODE	Specified as A0h (SPC-4).
SELECT REPORT	See Table 4-6.
ALLOCATION LENGTH	Refer to 3.1.
CONTROL	Refer to 3.3.

Table 4-6: REPORT LUNS SELECT REPORT CDB Field Notes

SELECT REPORT Field	Notes and References
00h	Shall be supported by returning all namespaces accessible to host software as individual LUNs as per LUN LIST field description – see section 6.6.
01h	Shall be supported by setting the LUN LIST LENGTH field to zero in the parameter data that is returned, indicating that there are no LUNs to report. This is not considered to be an error.
02h	Shall be supported by returning all namespaces accessible to host software as individual LUNs as per LUN LIST field description – see section 6.6.
10h	Shall be supported by setting the LUN LIST LENGTH field to zero in the parameter data that is returned, indicating that there are no LUNs to report. This is not considered to be an error.
11h	<p>If the LUN addressed in the command is LUN 0, Shall be supported by returning all namespaces accessible to host software as individual LUNs as per LUN LIST field description – see section 6.6.</p> <p>If the LUN addressed in the command is not LUN 0, Shall be supported by setting the LUN LIST LENGTH field to zero in the parameter data that</p>

	is returned, indicating that there are no LUNs to report. This is not considered to be an error.
12h	Shall be supported by setting the LUN LIST LENGTH field to zero in the parameter data that is returned, indicating that there are no LUNs to report. This is not considered to be an error.
All Others	Unspecified.

4.6 REQUEST SENSE (SPC-4)

REQUEST SENSE requests that sense data be returned to application client. Support requires returning sense data as defined in 6.5.

Table 4-7: REQUEST SENSE CDB Field Translations

REQUEST SENSE Field	Notes and References	
OPERATION CODE	Specified as 03h (SPC-4).	
DESC	CODE	Notes
	1b	Shall be supported by returning descriptor format sense data to application client. Refer to 6.5.2
	0b	Shall be supported by returning fixed format sense data to application client. Refer to 6.5.1.
ALLOCATION LENGTH	Refer to 3.1.	
CONTROL	Refer to 3.3.	

4.7 SECURITY PROTOCOL IN (SPC-4)

The SCSI SECURITY PROTOCOL IN command is used to retrieve information regarding security protocols and previous SECURITY PROTOCOL OUT commands. SECURITY PROTOCOL IN shall translate to an NVM Express Security Receive command.

Table 4-7: SECURITY PROTOCOL IN CDB Field Translations

SECURITY PROTOCOL IN Field	Notes and References	
OPERATION CODE	Specified as A2h (SBC-3).	
SECURITY PROTOCOL	Shall translate to Security Protocol field of NVM Express Security Receive command.	
SECURITY PROTOCOL SPECIFIC	Shall translate to SP Specific field of NVM Express Security Receive command.	
INC_512	CODE	Notes
	0b	Shall be supported. Indicates ALLOCATION LENGTH field specifies number of bytes to transfer.
	1b	Shall be supported. Indicates ALLOCATION LENGTH field specifies increments of 512 bytes.
ALLOCATION LENGTH	Shall translate to Allocation Length field of NVM Express Security Receive command.	
CONTROL	Refer to 3.3.	

4.8 SECURITY PROTOCOL OUT (SPC-4)

The SCSI SECURITY PROTOCOL OUT command is used to send data specifying operations to be performed. SECURITY PROTOCOL OUT command shall translate to an NVM Express Security Send command.

Table 4-8: SECURITY PROTOCOL OUT CDB Field Translations

SECURITY PROTOCOL OUT Field	Notes and References						
OPERATION CODE	Specified as B5h (SBC-3).						
SECURITY PROTOCOL	Shall translate to Security Protocol field of NVM Express Security Send command.						
SECURITY PROTOCOL SPECIFIC	Shall translate to SP Specific field of NVM Express Security Send command.						
INC_512	<table><tr><th>CODE</th><th>Notes</th></tr><tr><td>0b</td><td>Shall be supported. Indicates TRANSFER LENGTH field specifies number of bytes to transfer.</td></tr><tr><td>1b</td><td>Shall be supported. Indicates TRANSFER LENGTH field specifies increments of 512 bytes of data.</td></tr></table>	CODE	Notes	0b	Shall be supported. Indicates TRANSFER LENGTH field specifies number of bytes to transfer.	1b	Shall be supported. Indicates TRANSFER LENGTH field specifies increments of 512 bytes of data.
CODE	Notes						
0b	Shall be supported. Indicates TRANSFER LENGTH field specifies number of bytes to transfer.						
1b	Shall be supported. Indicates TRANSFER LENGTH field specifies increments of 512 bytes of data.						
TRANSFER LENGTH	Shall translate to Transfer Length field of NVM Express Security Send command.						
CONTROL	Refer to 3.3.						

4.9 SEND DIAGNOSTIC (SPC-4)

Since NVMe does not provide a method to request a self-test, this command shall be supported by taking no action and returning a GOOD status.

Table 4-9: SEND DIAGNOSTIC Field Translations

SEND DIAGNOSTIC Field	Notes and References
OPERATION CODE	Specified as 1Dh (SPC-4).
SELF-TEST CODE	Shall be set to 000b.
PF	Page Format, shall be set to 0b.
SELFTEST	May be set to 0b or 1b.
DEVOFFL	SCSI target device offline, unspecified.
UNITOFFL	Unit offline, unspecified.
PARAMETER LIST LENGTH	Shall be set to 0h.
CONTROL	Refer to 3.3.

4.10 START STOP UNIT (SPC-4)

START STOP UNIT requests that power states of the device transition to states specified by the command. Support for START STOP UNIT requires Set Features, Get Features and power state transitioning. Get Features and Set Features shall be used with namespace FFFFFFFFh.

Table 4-10: START STOP UNIT CDB Field Translations

START STOP UNIT Field	Notes and References
OPERATION CODE	Specified as 1Bh (SBC-3).
IMMED	Refer to 3.12.
POWER CONDITION MODIFIER	Refer to Table 4-11.
POWER CONDITION	Refer to Table 4-11.
NO_FLUSH	If set to 0b a Flush command shall precede any command that may result from a START STOP UNIT translation.
LOEJ	Unspecified. Indicates actions regarding loading or ejecting medium.
START	Refer to Table 4-11.
CONTROL	Refer to 3.3.

Table 4-11: Translation of Power States

POWER CONDITION	Name	POWER CONDITION MODIFIER	Notes	
00h	START_VALID	0h	START Code	Notes
			01h	A Set Features command shall be issued specifying power state 0
			00h	A Set Features command shall be issued specifying the lowest power state supported as indicated in the Get Features command.
01h	ACTIVE	0h	A Set Features command shall be issued specifying power state 0.	
02h	IDLE	0h	A Set Features command shall be issued specifying power state 1.	
		1h	A Set Features command shall be issued specifying power state 2.	
		2h	A Set Features command shall be issued specifying power state 3.	
03h	STANDBY	0h	A Set Features command shall be issued specifying power state N-2, where N is the last power state supported.	
		1h	A Set Features command shall be issued specifying power state N-1, where N is the last power state supported.	
07h	LU_CONTROL	0h	Unspecified.	

4.11 TEST UNIT READY (SPC-4)

TEST UNIT READY support requires a status to be returned to application client indicating the status of the logical unit. If the NVM Express device is ready to accept commands a status of GOOD shall be returned, otherwise a status of CHECK CONDITION and a sense key of NOT READY shall be returned.

Table 4-12: TEST UNIT READY CDB Field Translations

TEST UNIT READY Field	Notes and References
OPERATION CODE	Specified as 00h (SPC-4).
CONTROL	Refer to 3.3.

4.12 WRITE BUFFER (SPC-4)

The SCSI WRITE BUFFER command is used for testing and downloading of microcode and errors. Support for downloading of microcode requires the NVM Express Firmware Image Download and Firmware Activate commands.

Table 4-13: WRITE BUFFER CDB Field Translations

WRITE BUFFER Field	Notes and References	
OPERATION CODE	Specified as 3Bh (SPC-4).	
MODE		
	CODE	Notes
	05h	Download microcode, save, and activate mode. This mode shall be supported by issuing a Firmware Image Download command, followed by a Firmware Activate command.
	07h	Download microcode with offsets, save, and activate mode. If there is no knowledge of the total length of the image to be downloaded, the command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code. If the total length of the image to be downloaded in multiple pieces is understood/determinable, this mode shall be supported by issuing a Firmware Image Download command and, after the entire image has been downloaded, a subsequent Firmware Activate command.

	0Eh	Download microcode with offsets, save, and defer activate mode. This mode shall be supported by issuing a Firmware Image Download command.
	0Fh	Activate deferred microcode mode. This mode shall be supported by issuing a Firmware Activate command.
	All others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
BUFFER ID	Shall translate to Firmware Slot (FS) field of Firmware Activate command.	
BUFFER OFFSET	Shall translate to Offset (OFST) field of Firmware Image Download command.	
PARAMETER LIST LENGTH	Shall translate to Number of Dwords (NUMD) field of Firmware Image Download command. Refer to 3.9.	
CONTROL	Refer to 3.3.	

4.13 PERSISTENT RESERVE IN/OUT (SPC-4)

NVMe Reservation support was added in NVM-Express 1.1. A device based on this or a more recent specification identifies its support for these commands by setting IDENTIFY CONTROLLER field ONCS bit 5. If the NVMe device does not support this optional command, the SNTL shall terminate SCSI Reservation commands with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID OP CODE additional sense code.

Prior to translating SCSI commands for reservations, the host should have set the Host Identifier (HOSTID) feature if applicable. Refer to the NVMe specification.

Many NVMe reservation commands use a Reservation Type (RTYPE) field. This is translated from the SCSI RESERVATION TYPE value in Table 4-14, and is provided as reference for many of the Reservation translations.

Table 4-14: SCSI RESERVATION TYPE translation

SCSI RESERVATION TYPE	NVMe RESERVATION TYPE (RTYPE)
0h: No reservation	0h: Not a reservation holder
1h: WRITE EXCLUSIVE	1h
3h: EXCLUSIVE ACCESS	2h
5h: WRITE EXCLUSIVE-REGISTRANTS ONLY	3h

6h: EXCLUSIVE ACCESS – REGISTRANTS ONLY	4h
7h: WRITE EXCLUSIVE – ALL REGISTRANTS	5h
8h: EXCLUSIVE ACCESS- ALL REGISTRANTS	6h
All Others	Reserved: if encountered as part of a PERSISTENT RESERVE OUT, the SNTL shall terminate the command with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID FIELD IN PARAMETER LIST additional sense

4.13.1 PERSISTENT RESERVE IN

PERSISTENT RESERVE IN (SPC-4) is used to obtain information about persistent reservations and reservation keys (i.e., registrations) that are active within a device server. This command is used in conjunction with the PERSISTENT RESERVE OUT command (section 4.13.2).

The host shall support translating the SCSI CDB as defined in Table 4-15. See section 6.7 for details on filling the returned parameter list data for each Service Action.

Table 4-15: PERSISTENT RESERVE IN CDB Field Translations

PERSISTENT RESV IN Field	Notes and References	
OPERATION CODE	Specified as 5Eh (SPC-4).	
SERVICE ACTION		
	CODE	Notes
	00h	READ KEYS: Reads all registered reservation keys (i.e., registrations). This shall be translated using NVM Reservation Report command. Refer to section 6.7.1 for translation details.
	01h	READ RESERVATION: Reads the current persistent reservations. This shall be translated using NVM Reservation Report command. Refer to section 6.7.2 for translation details.

	02h	REPORT CAPABILITIES: Returns capability information. This shall be translated using Admin Identify Namespace and Get Feature commands. Refer to section 6.7.3 for translation details.
	03h	READ FULL STATUS: Reads complete information about all registrations and the persistent reservations, if any. This shall be translated using NVM Reservation Report command. Refer to section 6.7.4 for translation details.
	04h to 1Fh	Reserved. SNL shall terminate the command with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID FIELD IN CDB additional sense code.
ALLOCATION LENGTH	Refer to 3.1.	
CONTROL	Refer to 3.3.	

4.13.2 PERSISTENT RESERVE OUT

PERSISTENT RESERVE OUT (SPC-4) is used to request service actions that reserve a logical unit for the exclusive or shared use of a particular I_T nexus. The command uses other service actions to manage and remove such persistent reservations.

I_T nexuses performing PERSISTENT RESERVE OUT service actions are identified by a registered reservation key provided by the application client. An application client may use the PERSISTENT RESERVE IN command to obtain the reservation key, if any, for the I_T nexus holding a persistent reservation and may use the PERSISTENT RESERVE OUT command to preempt that persistent reservation.

The host shall support translating the SCSI CDB as defined in the following table. See section 6.7.5 for detailed translation for the command's Parameter List for each Service Action.

Table 4-16: PERSISTENT RESERVE OUT CDB

PERSISTENT RESERVE OUT Field	Notes and References	
OPERATION CODE	Specified as 5Fh (SPC-4).	
SERVICE ACTION		
	CODE	Notes

	00h	REGISTER: Register a reservation key with the device server or unregister a reservation key. This shall be translated into NVM Reservation Register command with IEKEY set to 0, RREGA set to either 000b (Register) or 001b (Unregister) depending on the service action (refer to 6.7.5).
	01h	RESERVE: Creates a reservation having a specified SCOPE and TYPE. This shall be translated into NVM Reservation Acquire command with IEKEY set to 0, and RACQA set to 000b (Acquire).
	02h	RELEASE: Releases the selected reservation. This shall be translated into NVM Reservation Release command with IEKEY set to 0, and RRELA set to 000b (Release)
	03h	CLEAR: Clears all reservation keys and reservations. This shall be translated into NVM Reservation Release command with IEKEY set to 0, and RRELA set to 001b (Clear)
	04h	PREEMPT: Preempts persistent reservations and/or removes registrations. This shall be translated into NVM Reservation Acquire command with IEKEY set to 0, and RACQA set to 001b (Preempt).
	05h	PREEMPT AND ABORT: Preempts persistent reservations and/or removes registrations and aborts all commands for preempted I_T Nexus. This shall be translated into NVM Reservation Acquire command with IEKEY set to 0, and RACQA set to 010b (Preempt and Abort).
	06h	REGISTER AND IGNORE EXISTING KEY: Register a reservation key with the device server or unregister a reservation key. This shall be translated into NVM Reservation Register command with IEKEY set to 1 and RREGA set to either 000b (Register) or 001b (Unregister) depending on the service action (refer to 6.7.5).
	07h	REGISTER AND MOVE: Register a reservation key for another I_T nexus with the device server and move a persistent reservation to that I_T nexus. This shall be translated into NVM Reservation Register command with IEKEY set to 0, and RREGA set to 010b (Replace).

	08h to 1Fh	Reserved. Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID FIELD IN CDB additional sense code.
TYPE	The host shall translate the SCSI TYPE to the NVM command's RTYPE field according to table 4-14 if used with NVM Reservation Acquire or Release command, and ignored otherwise.	
SCOPE	Only 00h is supported. If set to any other value, the SNTL shall terminate the command with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID FIELD IN CDB additional sense.	
PARAMETER LIST LENGTH	Refer to 3.9.	
CONTROL	Refer to 3.3.	

5 SCSI Block Command Mappings

This section defines the translation of SCSI command set extensions for direct-access block devices into NVM Express commands and structures.

5.1 COMPARE AND WRITE (SBC-3)

Support for COMPARE AND WRITE requires support for the NVM Express Compare and Write fused operation and is indicated in bit 0 of the Fused Operation Support field of the Identify Controller data structure.

Table 5-1: COMPARE AND WRITE CDB Field Translations

COMPARE AND WRITE Field	Notes and References
OPERATION CODE	Specified as 89h (SPC-4).
FUA_NV	Refer to 3.7.
FUA	Refer to 3.6.
DPO	Refer to 3.4.
WRPROTECT	Refer to 3.15.
LOGICAL BLOCK ADDRESS	Refer to 3.8.
GROUP NUMBER	Refer to 3.5.
NUMBER OF LOGICAL BLOCKS	Shall be translated to Number of Logical Blocks field of the NVM Express Compare command.
CONTROL	Refer to 3.3.

5.2 FORMAT UNIT (SBC-3)

Support for FORMAT UNIT requires the NVM Express Format NVM command and retention of NUMBER OF LOGICAL BLOCKS and LOGICAL BLOCK LENGTH mode parameter block descriptor fields as defined in 6.3.2.

Table 5-2: FORMAT UNIT CDB Field Translations

FORMAT UNIT Field	Notes and References	
OPERATION CODE	Specified as 04h (SPC-4).	
FMTPINFO	CODE	Notes
	0b	Shall be supported by formatting unit without protection.
	1b	May be supported if protection information is supported by formatting unit with protection information, refer to 5.2.1. Otherwise command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
LONGLIST	LONGLIST shall be ignored if FMTDATA field is set to zero. Otherwise:	
	CODE	Notes
	0b	Indicates a short parameter list header.
	1b	Indicates a long parameter list header.
FMTDATA	Unspecified. When set to 1b DEFECT LIST FORMAT is used.	
CPMLST	Unspecified. Indicates the type of defect list (complete list). Ignored when FMTDATA is set to 1.	
DEFECT LIST FORMAT	Unspecified. If FMTDATA is 1, specifies format of address descriptors in defect list.	
CONTROL	Refer to 3.3.	

5.2.1 FORMAT UNIT Parameter list

FORMAT UNIT Parameter list contains a header, an optional initialization pattern descriptor, and an optional defect list.

Table 5-3: FORMAT UNIT Parameter List Header (long and short) Translations

FORMAT UNIT Parameter List Header Field	Notes and References						
PROTECTION FIELD USAGE	Refer to 5.2.1.						
FOV	Unspecified						
DPRY	Unspecified						
DCRT	Unspecified						
STPF	Unspecified						
IP	Unspecified						
IMMED	Refer to 3.12.						
P_I_INFORMATION	SBC-3 specifies this as being set to zero.						
PROTECTION INTERVAL EXPONENT	<table border="1"> <thead> <tr> <th>Code</th><th>Notes</th></tr> </thead> <tbody> <tr> <td>0</td><td>Protection information interval shall equal the logical block length.</td></tr> <tr> <td>All others</td><td>Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.</td></tr> </tbody> </table>	Code	Notes	0	Protection information interval shall equal the logical block length.	All others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
Code	Notes						
0	Protection information interval shall equal the logical block length.						
All others	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.						
DEFECT LIST LENGTH	Unspecified						

5.2.1 Formatting with Protection Information

The type of protection to format a logical unit with is dependent upon the FTMPINFO field of FORMAT UNIT and the PROTECTION FIELD USAGE field of the FORMAT UNIT Parameter list.

Table 5-4: PROTECTION FIELD USAGE and FTMPINFO Translations

PROTECTION FIELD USAGE Code	FTMPINFO Code	Notes and References
000b	00b	Logical unit shall be formatted with type 0 protection. Protection Information field of Format NVM shall be set to 000b.
000b	10b	Logical unit shall be formatted with type 1 protection. Protection Information field of Format NVM shall be set to 001b.
000b	11b	Logical unit shall be formatted with type 2 protection. Protection Information field of Format NVM shall be set to 010b.
001b	11b	Logical unit shall be formatted with type 3 protection. Protection Information field of Format NVM shall be set to 011b.

All other values are unspecified.

5.3 READ(6), READ(10), READ(12), READ(16) (SBC-3)

Support for READ(6), READ(10), READ(12) or READ(16) requires an NVM Express Read command.

Table 5-5: READ CDB Field Notes

READ Field	Notes and References		
OPERATION CODE	Specified as (SBC-3):		
	SCSI Command	OPERATION CODE	
	READ(6)	08h	
	READ(10)	28h	
	READ(12)	A8h	
	READ(16)	88h	
FUA_NV*	Refer to 3.7.		
FUA*	Refer to 3.6.		
DPO*	Refer to 3.4.		
RDPROTECT*	Support for field requires setting PRACT and PRCHK of NVM Express read command:		
	RDPROTECT	PRACT	PRCHK
	000b	1	111b
	001b, 101b	0	111b
	010b	0	011b
	011b	0	000b
	100b	0	100b
	All other codes shall result in command termination with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.		
	LOGICAL BLOCK ADDRESS	Refer to 3.8.	
GROUP NUMBER*	Refer to 3.5.		
TRANSFER LENGTH	Refer to 3.14.		
CONTROL	Refer to 3.3.		
* READ(10), READ(12), and READ(16) only			

5.4 READ CAPACITY(10), READ CAPACITY(16) (SBC-3)

READ CAPACITY(10) and READ CAPACITY(16) requests that information regarding capacity and medium format be returned to application client. Support for READ CAPACITY(10) and READ CAPACITY(16) requires an Identify command and returning READ CAPACITY Parameter Data as specified in 6.4.

Table 5-6: READ CAPACITY CDB Field Translations

READ CAPACITY Field	Notes and References	
OPERATION CODE	Specified as (SBC-3):	
	SCSI Command	OPERATION CODE
	READ CAPACITY(10)	25h
	READ CAPACITY(16)	9Eh
SERVICE ACTION*	Set to 10h as specified in SPC-4.	
LOGICAL BLOCK ADDRESS	Shall be set to zero when PMI bit is not obsolete and set to 0b, otherwise command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.	
ALLOCATION LENGTH*	Refer to 3.1.	
PMI	This field is obsolete in some versions of SBC-3. If not obsolete, this field shall be translated:	
	CODE	Notes
	0b	Shall be supported by returning information on the last logical block.
	1b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
CONTROL	Refer to 3.3.	
* READ CAPACITY (16) only		

5.5 SYNCHRONIZE CACHE(10), SYNCHRONIZE CACHE(16) (SBC-3)

Support SYNCHRONIZE CACHE(10) and SYNCHRONIZE CACHE(16) requires an NVM Express Flush Command. All command specific fields are unspecified.

5.6 UNMAP (SBC-3)

Support for UNMAP requires NVM Express Dataset Management command.

Table 5-7: UNMAP CDB Field Translations

UNMAP Field	Notes and References	
OPERATION CODE	Specified as 42h (SBC-3).	
ANCHOR	CODE	Notes
	0b	Shall be supported by setting Attribute – Deallocate (AD) bit to 1b in NVM Express Dataset Management command
	1b	Command permissible for resource-provisioned devices.
GROUP NUMBER	Refer to 3.5.	
PARAMETER LIST LENGTH	Refer to 3.9.	
CONTROL	Refer to 3.3.	

5.6.1 UNMAP Parameter List

The parameter list specifies information regarding the block descriptors in Data-out buffer.

Table 5-8: UNMAP Parameter List

UNMAP Parameter List Field	Notes and References
UNMAP DATA LENGTH	Specifies number of bytes in Data-out buffer.
UNMAP BLOCK DESCRIPTOR DATA LENGTH	Specifies number of bytes of UNMAP block descriptors in Data-out buffer.

5.6.2 UNMAP Block Descriptor

The UNMAP Block Descriptor shall be translated into the Dataset Management Range Definition. NOTE: In SCSI, LBAs are stored in big endian, whereas NVMe are stored little endian.

Table 5-9: UNMAP Block Descriptor Field Notes

UNMAP Block Descriptor Field	Notes and References
UNMAP LOGICAL BLOCK ADDRESS	Shall translate to Starting LBA field
NUMBER OF LOGICAL BLOCKS	Shall translate to Length in Logical Blocks field.

5.7 VERIFY(10), VERIFY(12), VERIFY(16)(SBC-3)

Support for VERIFY(10), VERIFY(12) and VERIFY(16) requires a translation to the NVM Express Compare command.

Table 5-10: VERIFY CDB Field Translations

VERIFY	Notes and References	
OPERATION CODE	Specified as (SBC-3):	
	SCSI Command	OPERATION CODE
	VERIFY(10)	2Fh
	VERIFY (12)	AFh
	VERIFY (16)	8Fh
VRPROTECT	Refer to 3.17.	
DPO	Refer to 3.4.	
BYTCHK	Refer to 3.2 .	
LOGICAL BLOCK ADDRESS	Refer to 3.8.	
GROUP NUMBER	Refer to 3.5.	
VERIFICATION LENGTH	Refer to 3.16.	
CONTROL	Refer to 3.3.	

5.8 WRITE(6), WRITE(10), WRITE(12), WRITE(16) (SBC-3)

Support for WRITE(6), WRITE(10), WRITE(12), and WRITE(16) requires a translation to the NVM Express Write command.

Table 5-11: WRITE CDB Field Translations

WRITE	Notes and References	
OPERATION CODE	Specified as (SBC-3):	
	SCSI Command	OPERATION CODE
	WRITE(6)	0Ah
	WRITE(10)	2Ah
	WRITE(12)	AAh
	WRITE(16)	8Ah
FUA_NV*	Refer to 3.7.	
FUA*	Refer to 3.6.	
DPO*	Refer to 3.4.	
WRPROTECT*	Refer to 3.18.	
LOGICAL BLOCK ADDRESS	Refer to 3.8.	
GROUP NUMBER*	Refer to 3.5.	
TRANSFER LENGTH	Refer to 3.14.	
CONTROL	Refer to 3.3.	
* WRITE(10), WRITE(12), and WRITE(16) only		

5.9 WRITE LONG(10), WRITE LONG(16) (SBC-3)

WRITE LONG requests that a logical/physical block be marked as containing an error, or transfer data to the medium. Support requires the Write Uncorrectable command.

Table 5-12: WRITE LONG CDB Field Translations

WRITE LONG Field	Notes and References	
OPERATION CODE	Specified as (SBC-3):	
	SCSI Command	OPERATION CODE
	WRITE LONG(10)	3Fh
	WRITE LONG(16)	9Fh
SERVICE ACTION*	Shall be set to 11h as specified in SBC-3.	
COR_DIS		
	CODE	Notes
	1b	Shall be supported by issuing Write Uncorrectable command.
	0b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
WR_UNCOR		
	CODE	Notes
	1b	Shall be supported by issuing Write Uncorrectable command.
	0b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
PBLOCK		
	CODE	Notes
	0b	Shall be supported by issuing Write Uncorrectable command.
	1b	Command may be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and ILLEGAL FIELD IN CDB additional sense code.
LOGICAL BLOCK ADDRESS	Refer to 3.8.	
BYTE TRANSFER LENGTH	Unspecified.	
CONTROL	Refer to 3.3.	
* WRITE LONG(16) only		

6 SCSI Data

Some SCSI commands use Data-in and Data-out buffers to return information to the application client, or to specify additional command parameters. The content of Data-in and Data-out buffers when used in this manner is defined by SPC-4 and SBC-3. This section describes support for SCSI data.

6.1 INQUIRY Data

6.1.1 Standard INQUIRY Data Page

Table 6-1: Translation of Standard INQUIRY data

Standard INQUIRY data Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating support for PERIPHERAL DEVICE TYPE.
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device.
RMB	Shall be set to 0b indicating media cannot be removed.
VERSION	Shall be set to 06h indicating implementation of SPC-4.
NORMACA	Shall be set to 0b indicating an NACA bit set to 1 is unsupported and the ACA task attribute is unsupported.
HISUP	Shall be set to 1b indicating hierarchal addressing.
RESPONSE DATA FORMAT	Shall be set to 0010b indicating response data in SPC-4 format.
ADDITIONAL LENGTH	Shall be set to 1Fh indicating size of structure with no additional bytes.
SCCS	Shall be set to 0b indicating embedded storage arrays are unsupported.
ACC	Shall be set to 0b indicating access controls coordinator is unsupported.
TPGS	Shall be set to 0b indicating asymmetric LUN access is unsupported.
3PC	Shall be set to 0b indicating third party copy commands are unsupported.
PROTECT	Shall use Identify Namespace Data Structure End to End Data Protection Type Settings (DPS) field to determine value:

	Identify DPS Value	PROTECT Value	
	000b	0b	
	All others	1b	
ENCSEV	Shall be set to 0b indicating embedded enclosure services are unsupported.		
MULTIP	Shall be set using bit 0 of Multi-Interface Capabilities field of the Identify Controller Data Structure.		
ADDR16	Shall be set to 0b indicating 16-bit wide SCSI addresses (SCSI Parallel Interface Specific INQUIRY data) are unsupported.		
WBUS16	Shall be set to 0b indicating 16-bit wide data transfers (SCSI Parallel Interface Specific INQUIRY data) are unsupported.		
SYNC	Shall be set to 0b indicating synchronous data transfers (SCSI Parallel Interface Specific INQUIRY data) are unsupported.		
CMDQUE	Shall be set to 1b indicating support for Command Management Model.		
T10 VENDOR IDENTIFICATION	Refer to 3.12.		
PRODUCT IDENTIFICATION	Refer to 3.10.		
PRODUCT REVISION LEVEL	Refer to 3.11.		

6.1.2 Supported VPD Pages Data Page

Table 6-2: Translation of Supported VPD Pages VPD Page

Supported VPD Pages Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating support for PERIPHERAL DEVICE TYPE connected to logical unit.
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device.
PAGE LENGTH	Shall be set to indicate number of items supported VPD pages list requires.
Supported VPD Page List	List of supported pages.

6.1.3 Unit Serial Number VPD Page

The Unit Serial Number VPD page returns product serial number information to the application client. Support requires Identify Namespace Data Structure.

Table 6-3: Translation of Unit Serial Number VPD Page

Unit Serial Number VPD Page Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating support for PERIPHERAL DEVICE TYPE connected to logical unit.
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device.
PAGE CODE	Set to 80h as defined in SPC-4.
PAGE LENGTH	See section 6.1.3.1.
PRODUCT SERIAL NUMBER	See section 6.1.3.1.

6.1.3.1 PRODUCT SERIAL NUMBER and PAGE LENGTH fields

NVMe EUI64 Field	NVMe NGUID Field	Notes and References
non-zero	zero	See section 6.1.3.1.2.
zero	non-zero	See section 6.1.3.1.1.
non-zero	non-zero	See section 6.1.3.1.1.
non-zero	zero	See section 6.1.3.1.2.
zero	zero	See section 6.1.3.1.3. (Note: This case valid for NVMe 1.0 devices only.)

6.1.3.1.1 NGUID-based Translation

PAGE LENGTH shall be set to 40 indicating the page length in bytes.

PRODUCT SERIAL NUMBER shall be set to a 40 byte value by translating the NGUID field of the Identify Namespace Data Structure. The NGUID field shall be translated by converting each nibble into an ASCII equivalent representation, right aligning, and inserting a “_” after the 4th, 8th, 12th, 16th, 20th, 24th, 28th position, and a “.” after the 32nd position in the string. The first hexadecimal digit shall be the most significant four bits of the first byte (i.e., most significant byte) of the NGUID field. For example, 0x0123456789ABCDEF0123456789ABCDEF would be converted to “0123_4567_89AB_CDEF_0123_4567_89AB_CDEF.”

6.1.3.1.2 EUI64-based Translation

PAGE LENGTH shall be set to 20 indicating the page length in bytes.

PRODUCT SERIAL NUMBER shall be set to a 20 byte value by translating the IEEE Extended Unique Identifier (EUI64) field of the Identify Namespace Data Structure. The EUI64 field shall be translated by converting each nibble into an ASCII equivalent representation, right aligning, and inserting a “_” after

the 4th, 8th, 12th position, and a “.” after the 16th position in the string. For example, 0x0123456789ABCDEF would be converted to “0123_4567_89AB_CDEF.”

6.1.3.1.3 V1.0-compliant Devices

PAGE LENGTH should be set to 30 indicating the page length in bytes.

PRODUCT SERIAL NUMBER should be formed as follows:

- Bits 239:80 20 bytes of Serial Number (bytes 23:04 of Identify Controller data structure)
- Bits 79:72 ASCII representation of “_”
- Bits 71:08 ASCII representation of 32 bit Namespace Identifier (NSID)
- Bits 07:00 ASCII representation of “.”

The Namespace Identifier field shall be translated by converting each hexadecimal nibble into an ASCII equivalent representation and right aligning. For example, NSID 0x68430102 on a Controller with Serial Number “0123456789ABCDEF0123” would be converted to “0123456789ABCDEF_68430102.”

6.1.4 Device Identification Data Page

Table 6-4: Translation of Device Identification VPD Page

Device Identification VPD Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating device for PERIPHERAL DEVICE TYPE connected to logical unit
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device
PAGE CODE	Shall be set to 83h indicating Device Identification VPD Page
PAGE LENGTH	Shall be set to the size of the remaining bytes of Device Identification VPD Page.
Designation Descriptor List	<p>The designation descriptor list shall contain at least one designator in NAA IEEE Registered Extended designator format (refer to 6.1.4.1), T10 Vendor ID based designator format (refer to 6.1.4.3), SCSI Name String designator format (refer to 6.1.4.4) or EUI-64 designator format (refer to 6.1.4.5).</p> <p>It is strongly recommended that the designator descriptor list contain at least one designator in the EUI-64 designator format (refer to 6.1.4.5).</p>

6.1.4.1 NAA IEEE Registered Extended designator format

Table 6-5: NAA Designation Descriptor

NAA IEEE Registered Extended Designator Format Fields	Notes and References
PROTOCOL IDENTIFIER	Shall be set to 0h. PIV field shall indicate this field is reserved as no specific protocol to be identified.
CODE SET	Shall be set to 1h indicating associated fields are in binary format.
PIV	Shall be set to 0b indicating PROTOCOL IDENTIFIER field is reserved.
ASSOCIATION	Shall be set to 00b indicating DESIGNATOR field is associated with logical unit.
DESIGNATOR TYPE	Shall be set to 3h indicating NAA format and assignment authority.
DESIGNATOR LENGTH	Shall be set to the number of bytes of the DESIGNATOR field.
NAA	Shall be set to 6h indicating IEEE Registered NAA Extended.
IEEE COMPANY_ID	Shall be set using IEEE OUI Identifier field of the Identify Controller Data Structure (Note: refer to ECN 008).
VENDOR SPECIFIC IDENTIFIER	<p>VENDOR SPECIFIC IDENTIFIER and VENDOR SPECIFIC IDENTIFIER EXTENSION shall logically be combined to form a single 100-bit field.</p> <p>NVMe devices in which the EUI64 bit field in the Identify Namespace data structure contains a non-zero value shall form the 100-bit field by concatenating the EUI64 (64 bits) field followed by "0x000000000" (36 bits).</p> <p>For example, "0x0123456789ABCDEF" would be converted to "0x0123456789ABCDEF000000000".</p> <p>NVMe devices compliant with revision 1.0 that do not support the EUI64 field shall form the 100-bit field as:</p> <ul style="list-style-type: none"> • Bits 99:84: 16 bit PCI Vendor ID (bytes 01:00 of Identify Controller data structure) • Bits 83:32: Lower 52 bits of Serial Number (bytes 10:04 of Identify Controller data structure) • Bits 31:00: Namespace Identifier
VENDOR SPECIFIC IDENTIFIER EXTENSION	

	NVMe devices in which the EUI64 bit field is cleared to zero are not supported by this designator format.
--	---

6.1.4.2 NAA Locally Assigned Designator Format

Table 6-6: Translation of NAA Locally Assigned Designator

NAA Locally Assigned Designator Format Fields	Notes and References
PROTOCOL IDENTIFIER	Shall be set to 0h. PIV field shall indicate this field is reserved as no specific protocol to be identified.
CODE SET	Shall be set to 1h indicating associated fields are in binary format.
PIV	Shall be set to 0b indicating PROTOCOL IDENTIFIER field is reserved.
ASSOCIATION	Shall be set to 00b indicating DESIGNATOR field is associated with logical unit.
DESIGNATOR TYPE	Shall be set to 3h indicating NAA format and assignment authority.
DESIGNATOR LENGTH	Shall be set to 08h indicating length of designator.
NAA	Shall be set to 3h indicating Locally Assigned Designator.
LOCALLY ADMINISTERED VALUE	<p>For NVMe devices compliant with revision 1.1 or later, shall be set to the first 60 bits of the IEEE Extended Unique Identifier (EUI64) field of the Identify Namespace Data Structure.</p> <p>For NVMe devices compliant with revision 1.0, shall be set to the concatenation of the translation of PRODUCT IDENTIFICATION field from standard INQUIRY data as specified in section 3.10, and the following 60-bit field:</p> <ul style="list-style-type: none"> • Bits 59:44: 16 bit PCI Vendor ID (bytes 01:00 of Identify Controller data structure) • Bits 43:08: Lower 36 bits of Serial Number (bytes 08:04 of Identify Controller data structure) • Bits 07:00: Lower 8 bits of Namespace Identifier

6.1.4.3 T10 Vendor ID Based Descriptor

Table 6-7: T10 Vendor ID based Designation Descriptor

T10 Vendor ID Based Descriptor Field	Notes and References
PROTOCOL IDENTIFIER	Shall be set to 0h. PIV field shall indicate this field is reserved as no specific protocol to be identified.
CODE SET	Shall be set to 2h indicating associated fields are in ASCII format.
PIV	Shall be set to 0b indicating PROTOCOL IDENTIFIER field is reserved.
ASSOCIATION	Shall be set to 00b indicating DESIGNATOR field is associated with logical unit.
DESIGNATOR TYPE	Shall be set to 1h indicating T10 vendor ID based format.
DESIGNATOR LENGTH	Shall be set to size of T10 VENDOR IDENTIFICATION and VENDOR SPECIFIC IDENTIFIER fields.
T10 VENDOR IDENTIFICATION	Refer to 3.13.
VENDOR SPECIFIC IDENTIFIER	<p>For NVMe devices compliant with revision 1.1 or later, shall be set to the concatenation of the translation of PRODUCT IDENTIFICATION field from standard INQUIRY data as specified in 3.10, followed by one of the following fields (containing a non-zero value):</p> <ul style="list-style-type: none">the IEEE Extended Unique Identifier (EUI64) field of the Identify Namespace Data Structure; orthe NGUID field of the Identify Namespace Data Structure. <p>For NVMe devices compliant with revision 1.0, shall be set to the concatenation of the translation of PRODUCT IDENTIFICATION field from standard INQUIRY data as specified in 3.10, and the following 100-bit field:</p> <ul style="list-style-type: none">Bits 99:84: 16 bit PCI Vendor ID (bytes 01:00 of Identify Controller data structure)Bits 83:32: Lower 52 bits of Serial Number (bytes 10:04 of Identify Controller data structure)Bits 31:00: Namespace Identifier

6.1.4.4 SCSI Name String Designator

Table 6-8: SCSI name string Designation Descriptor

SCSI Name String Descriptor Field	Notes and References
PROTOCOL IDENTIFIER	Shall be set to 0h. PIV field shall indicate this field is reserved as no specific protocol to be identified.
CODE SET	Shall be set to 3h indicating associated fields are in UTF-8 format.
PIV	Shall be set to 0b indicating PROTOCOL IDENTIFIER field is reserved.
ASSOCIATION	Shall be set to 00b indicating DESIGNATOR field is associated with logical unit.
DESIGNATOR TYPE	Shall be set to 8h indicating SCSI name string format and assignment authority.
DESIGNATOR LENGTH	Shall be set to size of SCSI NAME STRING field.
SCSI NAME STRING	See section 6.1.4.4.1.

6.1.4.4.1 SCSI NAME STRING field

NVMe EUI64 Field	NVMe NGUID Field	Notes and References
non-zero	zero	See section 6.1.4.4.1.2.
zero	non-zero	See section 6.1.4.4.1.1.
non-zero	non-zero	Section 6.1.4.4.1.1 or 6.1.4.4.1.2 shall be followed. Two separate descriptors, one based on each of the two sections above, may be returned.
non-zero	zero	See section 6.1.4.4.1.2.
zero	zero	See section 6.1.4.4.1.3. (Note: This case valid for NVMe 1.0 devices only.)

6.1.4.4.1.1 NGUID-based Translation

Shall be set to a 36 byte UTF-8 character field comprised of the four UTF-8 characters 'eui.' concatenated with UTF-8 representation of the 32 hexadecimal digits corresponding to the 128 bit NGUID field of the Identify Namespace Data Structure. The first hexadecimal digit shall be the most significant four bits of the first byte (i.e., most significant byte) of the NGUID field.

6.1.4.4.1.2 EUI64-based Translation

Shall be set to a 20 byte UTF-8 character field comprised of the four UTF-8 characters 'eui.' concatenated with UTF-8 representation of the 16 hexadecimal digits corresponding to the 64 bit EUI64 field of the Identify Namespace Data Structure. The first hexadecimal digit shall be the most significant four bits of the first byte (i.e., most significant byte) of the EUI-64 field.

6.1.4.4.1.3 V1.0-compliant Devices

Shall be set to a 68 byte UTF-8 character field comprised of 4 bytes of UTF-8 representation of 2 byte PCI Vendor ID, plus 40 bytes of Model Number, plus 4 bytes of UTF-8 representation of Namespace ID, plus 20 bytes of Serial Number. Note: the start of the string (MSB) is at the lowest byte offset and the end of the string (LSB) is at the highest byte offset:

- Bytes 67:48: 20 bytes of Serial Number(bytes 23:04 of Identify Controller data structure)
- Bytes 47:44: 4 bytes of Namespace ID (UTF-8 representation)
- Bytes 43:04: 40 bytes of Model Number(bytes 63:24 of Identify Controller data structure)
- Bytes 03:00: 4 bytes of PCI Vendor ID (UTF-8 representation) (bytes 01:00 of Identify Controller converted to 4 UTF-8 characters)

6.1.4.5 EUI-64 based Designator

Table 6-9: EUI-64 Designation Descriptor

EUI64 Descriptor Field	Notes and References
PROTOCOL IDENTIFIER	Shall be set to 0h. PIV field shall indicate this field is reserved as no specific protocol to be identified.
CODE SET	Shall be set to 1h indicating associated fields are in binary format.
PIV	Shall be set to 0b indicating PROTOCOL IDENTIFIER field is reserved.
ASSOCIATION	Shall be set to 00b indicating DESIGNATOR field is associated with logical unit.
DESIGNATOR TYPE	Shall be set to 2h indicating EUI-64 based identifier.
DESIGNATOR LENGTH	Shall be set to 08h if the EUI-64 field is populated from the NVMe EUI64 field. Shall be set to 10h if the EUI-64 field is populated from the NVMe NGUID field.
EUI-64 Field	See section 6.1.4.5.1.

6.1.4.5.1 EUI-64 Field

NVMe EUI64 Field	NVMe NGUID Field	Notes and References
non-zero	zero	Shall contain the value of the NVMe EUI64 field.
zero	non-zero	Shall contain the value of the NVMe NGUID field.
non-zero	non-zero	Shall contain the value of either the NVMe EUI64 field or the NVMe NGUID field. Two separate descriptors, one based on each of the two fields above respectively, may be returned.
non-zero	zero	Shall contain the value of the NVMe EUI64 field.
zero	zero	Not supported – the EUI-64 Designator descriptor format shall not be used.

6.1.5 Extended INQUIRY Data VPD Page

This page is optional.

Table 6-10: Translation of Extended INQUIRY Data VPD Page

Extended INQUIRY Data Field	Notes and References																		
PERIPHERAL QUALIFIER	Shall be set to 000b indicating device for PERIPHERAL DEVICE TYPE connected to logical unit																		
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device																		
PAGE CODE	Shall be set to 86h indicating Extending INQUIRY Data VPD																		
PAGE LENGTH	Set to 3Ch																		
ACTIVATE MICROCODE	Shall be set to 10b indicating microcode will be activated after a hard reset.																		
SPT	<p>SPT shall be translated using the DPC field of the Identify Namespace Data Structure:</p> <table border="1"> <thead> <tr> <th>DPC Value (bits 0-2)</th><th>SPT Value</th></tr> </thead> <tbody> <tr> <td>000b</td><td><i>Undefined</i></td></tr> <tr> <td>001b</td><td>000b</td></tr> <tr> <td>010b</td><td>010b</td></tr> <tr> <td>011b</td><td>001b</td></tr> <tr> <td>100b</td><td>100b</td></tr> <tr> <td>101b</td><td>011b</td></tr> <tr> <td>110b</td><td>101b</td></tr> <tr> <td>111b</td><td>111b</td></tr> </tbody> </table>	DPC Value (bits 0-2)	SPT Value	000b	<i>Undefined</i>	001b	000b	010b	010b	011b	001b	100b	100b	101b	011b	110b	101b	111b	111b
DPC Value (bits 0-2)	SPT Value																		
000b	<i>Undefined</i>																		
001b	000b																		
010b	010b																		
011b	001b																		
100b	100b																		
101b	011b																		
110b	101b																		
111b	111b																		
GRD_CHK	If DPS field of Identify Namespace Data Structure is 000b, this field shall be set to 0b, otherwise set to 1b.																		
APP_CHK	If DPS field of Identify Namespace Data Structure is 000b, this field shall be set to 0b, otherwise set to 1b.																		
REF_CHK	If DPS field of Identify Namespace Data Structure is 000b, this field shall be set to 0b, otherwise set to 1b.																		
UASK_SUP	Shall be set to 1b indicating sense key specific data is returned for UNIT ATTENTION sense key.																		
GROUP_SUP	Shall be set to 0b indicating grouping function is unsupported.																		
PRIOR_SUP	Shall be set to 0b indicating command priority is unsupported.																		
HEADSUP	Shall be set to 0b indicating HEAD OF QUEUE task attribute is unsupported.																		

ORDSUP	Shall be set to 0b indicating ORDERED task attribute is unsupported.
SIMPSUP	Shall be set to 0b indicating SIMPLE task attribute is unsupported.
WU_SUP	Shall be set to 0b indicating WR_UNCOR field of WRITE LONG is unsupported if the NVMe Write Uncorrectable command is unsupported. Shall be set to 1b indicating WR_UNCOR field of WRITE LONG is supported if the NVMe Write Uncorrectable command is supported.
CRD_SUP	Shall be set to 1b indicating COR_DIS field of WRITE LONG is supported if the NVMe Write Uncorrectable command is supported. Otherwise shall be set to 0b.
NV_SUP	Shall be set to 0b indicating a non-volatile cache is unsupported, i.e. FUA_NV is unsupported.
V_SUP	Shall be set using the Volatile Write Cache (VMC) field of the Identify Controller Data Structure.
P_I_I_SUP	Shall be set to 0b indicating protection information intervals are unsupported.
LUICLR	Shall be set to 1b indicating unit attentions are cleared according to SPC-4.
R_SUP	Shall be set to 0b indicating referrals are unsupported.
CBCS	Shall be set to 0b indicating capability-based command security is unsupported.
MULTI I_T NEXUS MICROCODE DOWNLOAD	Shall be set to 0h indicating handling of microcode downloads is vendor specific.
EXTENDED SELF-TEST COMPLETION MINUTES	Shall be set to 0h indicating that this field is unsupported.
POA_SUP	Shall be set to 0b indicating power on activation is unsupported.
HRA_SUP	Shall be set to 0b indicating hard reset activation is unsupported.
VSA_SUP	Shall be set to 0b indicating vendor specific activation is unsupported.

MAXIMUM SUPPORTED SENSE DATA LENGTH	Shall be set to 0h indicating that the device does not report a maximum sense data length.
--	---

6.1.6 Block Limits VPD Page

Table 6-11: Translation of Block Limits VPD Page

Block Limits VPD Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating device for PERIPHERAL DEVICE TYPE connected to logical unit
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device
PAGE CODE	Shall be set to B0h indicating Block Limits VPD Page
PAGE LENGTH	Shall be set to 003Ch.
WSNZ	Unspecified
MAXIMUM COMPARE AND WRITE LENGTH	Shall be set to 00h if Fused Operation is not supported; May be set to a non-zero value that is less than or equal to the value in MAXIMUM TRANSFER LENGTH field if Fused Operation is supported.
OPTIMAL TRANSFER LENGTH GRANULARITY	Unspecified
MAXIMUM TRANSFER LENGTH	Shall be set to value calculated according to method described in NVMe v1.1 Identify Controller Data Structure: Maximum Data Transfer Size (MDTS)
OPTIMAL TRANSFER LENGTH	Unspecified
MAXIMUM PREFETCH LENGTH	Unspecified
MAXIMUM UNMAP LBA COUNT	Shall be set to 0000_0000h if Dataset Management command – Deallocate (AD) attribute is not supported. Shall be set to non-zero value if Dataset Management command – Deallocate (AD) attribute is supported.
MAXIMUM UNMAP BLOCK DESCRIPTOR COUNT	Shall be set to 0000_0000h if Dataset Management command – Deallocate (AD) attribute is not supported. Shall be set to 0000_0100h if Dataset Management command – Deallocate (AD) attribute is supported.
OPTIMAL UNMAP GRANULARITY	Unspecified
UNMAP GRANULARITY ALIGNMENT	Unspecified
UGAVALID	Unspecified

MAXIMUM WRITE SAME LENGTH	Unspecified
---------------------------	-------------

6.1.7 Block Device Characteristics VPD Page

Table 6-12: Translation of Block Device Characteristics VPD Page

Block Device Characteristics VPD Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating device for PERIPHERAL DEVICE TYPE connected to logical unit
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device
PAGE CODE	Shall be set to B1h indicating Device Identification VPD Page
PAGE LENGTH	Shall be set to 3Ch.
MEDIUM ROTATION RATE	Shall be set to 0001h indicating a non-rotating device (SSD).
NOMINAL FORM FACTOR	Shall be set to 0h indicating form factor not reported.

6.1.8 Logical Block Provisioning VPD Page

Table 6-13: Translation of Block Device Characteristics VPD Page

Block Limits VPD Field	Notes and References
PERIPHERAL QUALIFIER	Shall be set to 000b indicating device for PERIPHERAL DEVICE TYPE connected to logical unit
PERIPHERAL DEVICE TYPE	Shall be set to 00h indicating direct access block device
PAGE CODE	Shall be set to B2h indicating Block Limits VPD Page
PAGE LENGTH	Shall be set to 04h.
THRESHOLD EXPONENT	Shall be set to 00h to indicate that there are no thin provisioning thresholds This would require modification if thin-provisioning is supported
LBPV	Shall be set to 0 if Dataset Management command – Deallocate (AD) attribute is not supported. Shall be set to 1 if Dataset Management command – Deallocate (AD) attribute is supported. Shall be set to 1 if PROVISIONING TYPE is set to 1 or 2. This is reporting whether use of UNMAP to unmap LBAs is supported.
LBPWS	Shall be set to 0, to indicate that use of WRITE SAME(16) to unmap LBAs is not supported.
LBPWS10	Shall be set to 0, to indicate that use of WRITE SAME(10) to unmap LBAs is not supported.
LBPRZ	Shall be set to 1 if Dataset Management command – Deallocate (AD) attribute is supported and the device returns all zeros for reads of deallocated LBAs, otherwise set to 0.

ANC_SUP	<p>Shall be set to 0, to indicate that setting the ANCHOR bit in UNMAP is not supported if the namespace is not resource or thin provisioned.</p> <p>Shall be set to 1, to indicate that setting the ANCHOR bit in UNMAP is supported if the namespace is resource provisioned. May be set to 1 if the namespace is thin provisioned.</p>
DP	<p>Shall be set to 0, to indicate that no Provisioning Group Descriptors follow.</p> <p>This would require modification if thin-provisioning is supported</p>
PROVISIONING TYPE	<p>Shall be set to 0 (Full) if Dataset Management command – Deallocate (AD) attribute is not supported and Identify Namespace Data = NSFEAT bit 0 is reported “0” indicating that the namespace is not thin-provisioned.</p> <p>Shall be set to 1(Resource) if Dataset Management command – Deallocate (AD) attribute is supported and Identify Namespace Data – NSFEAT bit 0 is reported “0”, indicating that the namespace is resource-provisioned.</p> <p>Shall be set to 2(Thin) if Dataset Management command – Deallocate (AD) attribute is supported and Identify Namespace Data – NSFEAT bit 0 is reported “1”, indicating that the namespace is thin-provisioned.</p>

6.2 Log Pages

6.2.1 Supported Log Pages

Table 6-14: Translation of Supported Log Pages

Supported Log Pages Log Page Field	Notes and References	
PAGE CODE	Set to 00h as specified in SPC-4.	
SPF	Set to 0b as specified in SPC-4.	
DS	Set to 0b as specified in SPC-4.	
SUBPAGE CODE	Set to 00h as specified in SPC-4.	
PAGE LENGTH	Shall be set to 03h or 04h indicating the length of Supported Pages List.	
Supported Pages List	PAGE CODE	Notes
	00h	Indicates support for Supported Log Pages log page. (Mandatory)
	0Dh	Indicates support for Temperature Log Page log page. (Mandatory)
	11h	Indicates support for Solid State Media log page. (Optional)
	2Fh	Indicates support for Informational Exceptions log page. (Mandatory)

6.2.2 Temperature Log Page

Table 6-15: Translation of Temperature Log Page

Temperature Log Page Field	Notes and References
DS	Unspecified.
SPF	Shall be set to 0b as specified in SPC-4.
PAGE CODE	Shall be set to 0Dh as specified in SPC-4.
SUBPAGE CODE	Shall be set to 00h as specified by SPC-4.
PAGE LENGTH	Shall be set to 0Ch as specified by SPC-4.
PARAMETER CODE	Shall be set to 0000h as specified in SPC-4.
FORMAT AND LINKING	Shall be set to 11b indicating parameter is in binary format.
TMC	Shall be set to 00b for LOG SENSE as specified by SPC-4 or ignored.
ETC	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
TSD	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
DU	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
PARAMETER LENGTH	Shall be set to 02h as specified in SPC-4.

TEMPERATURE	Shall be set to Temperature field of Get Log Page – SMART / Health Information Log. This shall require a conversion from Kelvin to Celsius.
PARAMETER CODE	Shall be set to 0001h as specified in SPC-4.
FORMAT AND LINKING	Shall be set to 11b indicating parameter is in binary format.
TMC	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
ETC	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
TSD	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
DU	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
PARAMETER LENGTH	Shall be set to 02h as specified in SPC-4.
REFERENCE TEMPERATURE	Shall be set using Temperature Threshold Feature of Get Features command. This shall require a conversion from Kelvin to Celsius

6.2.3 Solid State Media Log Page

Table 6-16: Translation of Solid State Media Log Page

Temperature Log Page Field	Notes and References
DS	Unspecified.
SPF	Shall be set to 0b as specified in SPC-4.
PAGE CODE	Shall be set to 11h as specified in SPB-3.
SUBPAGE CODE	Shall be set to 00h as specified by SPB-3.
PAGE LENGTH	Should be set to 08h as specified by SPB-3.
PARAMETER CODE	Shall be set to 0001h as specified in SPB-3.
FORMAT AND LINKING	Shall be set to 11b indicating parameter is in binary format.
TMC	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
ETC	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
TSD	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
DU	Shall be set to 0 for LOG SENSE as specified by SPC-4 or ignored.
PARAMETER LENGTH	Shall be set to 04h as specified in SBC-3.
PERCENTAGE USED ENDURANCE INDICATOR	<ol style="list-style-type: none"> 1. <u>Issue Get Features command to retrieve the SMART / Health Information Log.</u> 2. <u>Set the PERCENTAGE USED ENDURANCE INDICATOR field in the Solid State Media log page from the Percentage Used field in the Get Features result.</u>

6.2.4 Information Exceptions Log Page

Table 6-17: Translation of Informational Exceptions Log Page

Informational Exceptions Log Page Field	Notes and References
PAGE CODE	Set to 2Fh as specified in SPC-4.
SPF	Set to 0b as specified in SPC-4.
DS	Set to 0b as specified in SPC-4.
SUBPAGE CODE	Set to 00h as specified in SPC-4.
PAGE LENGTH	Shall be set to 8h indicating the length of remaining log page.
PARAMETER CODE	Set to 0000h as specified by SPC-4.
FORMAT AND LINKING	Shall be set to 11b indicating parameter is a binary format list parameter.
TMC	Set to 00b as specified in SPC-4.
ETC	Set to 0b as specified in SPC-4.
TSD	Set to 1b indicating log parameter disabled.
DU	Set to 0b as specified in SPC-4.
PARAMETER LENGTH	Shall be set to 04h indicating 4 parameters.
INFORMATIONAL EXCEPTION ADDITIONAL SENSE CODE	Shall be set to 0h.
INFORMATIONAL EXCEPTION ADDITIONAL SENSE CODE QUALIFIER	Shall be set to 0h.
MOST RECENT TEMPERATURE READING	Shall be supported using Temperature field of Get Log Page SMART / Health Information Log. A conversion to Celsius from Kelvin must occur.

6.3 Mode Parameters

Mode Parameters specify information regarding a logical unit. MODE SELECT is used by application clients to set parameter values. MODE SENSE is used to return current values of parameter data to application client. An NVM Express driver may support a mode parameter by mapping the value to an NVM Express data structure if possible or by storing the parameter value internally within the driver. There is no distinction between current and saved mode pages. The mode parameter list contains a mode parameter header, zero or more block descriptors, followed by zero or more variable-length mode pages.

6.3.1 Mode Parameter Headers

Mode Parameter Header(6) is used by MODE SELECT(6) and MODE SENSE(6) commands. Mode Parameter Header(10) is used by MODE SELECT(10) and MODE SENSE(10). Mode Parameter Header(6) and Mode Parameter Header(10) have similar fields and are translated in Table 6-18. When not specified in the description, each field is applicable to both MODE SENSE and MODE SELECT commands.

Table 6-18: Mode Parameter Header Field Translations

Mode Parameter Header Field	Notes and References	
MODE DATA LENGTH	Shall indicate number of bytes of data to transfer. The mode data length does not include the number of bytes in the MODE DATA LENGTH field. When translating MODE SELECT commands, this field is reserved.	
MEDIUM TYPE	Shall be set to 00h indicating a direct access block device.	
WP	Shall be set to 1b if bit 3 of the Critical Warning field of the SMART / Health Information Log page is set to 1b (i.e. media has been placed in read only mode.). Otherwise shall be set to 0b.	
DPOFUA	When used with MODE SELECT this field is reserved. Shall be set to 1b indicating DPO and FUA are supported.	
LONGLBA*	This field shall be ignored for MODE SELECT 10.	
	LONGLBA	Notes
	0b	Shall be set to 0b indicating mode parameter block descriptors are 8 bytes long should there be any.
	1b	Shall be set to 1b indicating mode parameter block descriptor(s) are 16 bytes long should there be any.
BLOCK DESCRIPTOR LENGTH	The value of this field is determined by the LONGLBA field and shall indicate the number of bytes of all block descriptors:	
	LONGLBA	Notes
	0b	Shall indicate the number if block descriptors multiplied by 8.

	1b	Shall indicate the number of block descriptors multiplied by 16.
	Note that this length does not include mode pages that may follow the last block descriptor should there be any.	
* Mode Parameter Header 10 only.		

6.3.2 Mode Parameter Block Descriptor

Mode Parameter Block Descriptors may be optionally supported and are used to specify the size and number of blocks of a logical unit. Table 6-19 describes how to determine the initial values of the parameters. Support for this parameter when used along with a FORMAT UNIT command requires retention of the fields within the driver during a MODE SELECT command.

Table 6-19: Mode Parameter Block Descriptor Field Translations

Mode Parameter Block Descriptor Field	Notes and References
NUMBER OF LOGICAL BLOCKS	Shall be determined through translation of Namespace Capacity of Identify Namespace Data Structure.
LOGICAL BLOCK LENGTH	Shall be determined through translation of Formatted LBA Size (FLBAS) field of Identify Namespace Data Structure.

6.3.3 Mode Page Formats

The following sections define the translation for the supported mode pages.

6.3.3.1 Read-Write Error Recovery Mode Page

Table 6-20: Read-Write Error Recovery Mode Page Field Translations

Read-Write Recovery Mode Page Field	Notes and References
PS	Shall be set to 0b indicating device server is unable to save parameters.
SPF	Shall be set to 0b indicating page 0 mode page format is being used
PAGE CODE	Shall be set to 01h indicating Read-Write Error Recovery Mode Page.
PAGE LENGTH	Shall be set to 0Ah indicating number of bytes of mode parameters.
DCR	Shall be set to 0b indicating the use of additional information (e.g., ECC bytes) for data error recovery is allowed.
DTE	Shall be set to 0b indicating that upon detection of a recovered error, the device shall not terminate the data transfer to the Data-In Buffer for a read command or the data transfer to the Data-Out Buffer for a write command.
PER	Shall be set to 0b
EER	Unspecified.
RC	Shall be set to 0b, indicating that error recovery operations that cause delays during the data transfer are acceptable; data shall not be fabricated.
TB	Unspecified.
ARRE	Shall be set to 1b, indicating that the device shall enable automatic read reassignment for LBAs referencing logical blocks for which a recovered error occurs during a read medium operation.
AWRE	Shall be set to 1b, indicating that the device shall enable automatic write reassignment for LBAs referencing logical blocks for which a recovered or unrecovered error occurs during a write medium operation.
READ RETRY COUNT	Unspecified.

WRITE RETRY COUNT	Unspecified.				
RECOVERY TIME LIMIT	<p>The value of this field is dependent on the command:</p> <table> <tr> <td>MODE SENSE</td><td>Shall be set to the value of the Time Limited Error Recovery field of the NVM Express Get Features command with Error Recovery Identifier.</td></tr> <tr> <td>MODE SELECT</td><td>The value of the Time Limited Error Recovery field of the NVMe Set Features Cmd with Error Recovery Feature Identifier shall be set to this nearest rounded-up 100 millisecond value.</td></tr> </table> <p>The value shall be converted to 100 millisecond units when sending to the device (MODE SELECT) and the value shall be converted to 1 millisecond units when receiving from the device (MODE SENSE).</p>	MODE SENSE	Shall be set to the value of the Time Limited Error Recovery field of the NVM Express Get Features command with Error Recovery Identifier.	MODE SELECT	The value of the Time Limited Error Recovery field of the NVMe Set Features Cmd with Error Recovery Feature Identifier shall be set to this nearest rounded-up 100 millisecond value.
MODE SENSE	Shall be set to the value of the Time Limited Error Recovery field of the NVM Express Get Features command with Error Recovery Identifier.				
MODE SELECT	The value of the Time Limited Error Recovery field of the NVMe Set Features Cmd with Error Recovery Feature Identifier shall be set to this nearest rounded-up 100 millisecond value.				

6.3.3.2 Caching Mode Page

Table 6-21: Control Mode Page Field Translations

Caching Mode Page Field	Notes and References				
Write Back Cache Enable (WCE)	<p>The value of this field is dependent on the command:</p> <table> <tr> <td>MODE SENSE</td><td>Shall be set to the value of Volatile Write Cache Enable field of the NVM Express Get Features command with Volatile Write Cache Feature Identifier.</td></tr> <tr> <td>MODE SELECT</td><td>Shall be set to the value of Volatile Write Cache Enable field of the NVM Express Set Features command with the Volatile Write Cache Feature Identifier to be sent to the NVMe device.</td></tr> </table>	MODE SENSE	Shall be set to the value of Volatile Write Cache Enable field of the NVM Express Get Features command with Volatile Write Cache Feature Identifier.	MODE SELECT	Shall be set to the value of Volatile Write Cache Enable field of the NVM Express Set Features command with the Volatile Write Cache Feature Identifier to be sent to the NVMe device.
MODE SENSE	Shall be set to the value of Volatile Write Cache Enable field of the NVM Express Get Features command with Volatile Write Cache Feature Identifier.				
MODE SELECT	Shall be set to the value of Volatile Write Cache Enable field of the NVM Express Set Features command with the Volatile Write Cache Feature Identifier to be sent to the NVMe device.				
All other fields shall be unspecified.					

6.3.3.3 Control Mode Page

Table 6-22: Control Mode Page Field Translations

Control Mode Page Field	Notes and References
PS	Shall be set to 0b indicating device server is unable to save parameters.
SPF	Shall be set to 0b indicating page_0 mode page format is being used.
PAGE CODE	Shall be set to 0Ah indicating Control Mode Page.
PAGE LENGTH	Shall be set to 0Ah indicating number of bytes of mode parameters.
TST	Shall be set to 000b indicating one task set for all I_T nexuses.
TMF_ONLY	Shall be set to 0b. Note that ACA is not supported.
DPICZ	Shall be set to 0b to indicate that checking of protection information bytes is enabled. Otherwise shall be set to 1b to indicate that checking of protection information is disabled on commands with the RDPROTECT, VRPROTECT or ORPROTECT field set to zero.
D_SENSE	Shall be set to 1b indicating that sense data is returned in descriptor format.
GLTSD	Shall be set to 1b indicating that the logical unit does not implicitly save log parameters.
RLEC	Shall be set to zero indicating log exception conditions are not reported.
QUEUE ALGORITHM MODIFIER	Shall be set to one indicating commands may be reordered.
NUAR	Shall be set to 0b. Note that reservations are not supported.
QERR	Shall be set to 01b.
RAC	Shall be set to 0b indicating the device server may return BUSY status for any length of time while BUSY condition persists.

UA_INTLCK_CTRL	Shall be set to 00b indicating that unit attention condition is cleared in the same transaction as a CHECK CONDITION status.
SWP	Shall be set to 0b indicating no support for software write protect.
ATO	Shall be set to 1b indicating that: <ul style="list-style-type: none"> a) the LOGICAL BLOCK APPLICATION TAG shall not be modified when the NVMe namespace is formatted with any type of protection information; and b) the LOGICAL BLOCK REFERENCE TAG shall not be modified when the NVMe namespace is formatted with type 3 protection information.
TAS	Shall be set to one indicating that aborted commands shall be completed with TASK ABORTED status.
AUTOLOAD MODE	Shall be set to 000b indicating medium is loaded for full access when inserted.
BUSY TIMEOUT PERIOD	Shall be set to FFFFh indicating unlimited busy timeout.
EXTENDED SELF-TEST COMPLETION TIME	Shall be set to 0000h, SMART self-tests are unsupported.

6.3.3.4 Power Condition Control Mode Page

Table 6-23: Translation of Power Condition Control Mode Page

Power Condition Control Mode Page Field	Notes and References
PS	Shall be set to 0h indicating page not savable.
SPF	Shall be set to 0h indicating page_0 page format is used.
PAGE CODE	Set to 1Ah as specified by SPC-4.
PAGE LENGTH	Set to 26h as specified by SPC-4. Indicates length of page.
PM_BG_PRECEDENCE	Shall be set to 00H, indicating vender specific power management and background interactions.
STANDBY_Y	Timers are not supported in NVM Express. When processing a MODE SENSE command, these fields shall be returned as zero. When processing a MODE SELECT command, the command may be terminated with CHECK CONDITION status, sense key set to ILLEGAL REQUEST and additional sense code set to INVALID FIELD IN PARAMETER LIST.
IDLE_C	
IDLE_B	
IDLE_A	
STANDBY_Z	
IDLE_A CONDITION TIMER	
STANDBY_Z CONDITION TIMER	
IDLE_B CONDITION TIMER	
IDLE_C CONDITION TIMER	
STANDBY_Y CONDITION TIMER	

6.3.3.5 Return all pages Mode Page

Table 6-24: Mode pages returned

Mode Pages (returned in ascending numerical order)
0x01
0x08
0x0A
0x1A
0x1C

6.3.3.6 Return all pages and subpages Mode Page

Table 6-25: Mode pages returned

Mode Pages (returned in ascending numerical order)	SubPage
0x01	0x00
0x08	0x00
0x0A	0x00
0x1A	0x00
0x1C	0x00

6.3.3.7 Informational Exceptions Control Mode Page

Table 6-26: Translation of Informational Exceptions Control Mode Page

Informational Exceptions Control Mode Page Field	Notes and References		
PS	Shall be set to 0b indicating page not savable.		
SPF	Shall be set to 0b indicating page_0 mode page format is used.		
PAGE CODE	Set to 1Ch as specified by SPC-4.		
PAGE LENGTH	Shall be set to 0Ah indicating length of page.		
LOGERR	Shall be set to 0b indicating device server may log informational exception conditions in the Informational Exceptions log page.		
EBACKERR	Shall be set to 0b indicating reporting of informational exception conditions occurring during processing of background self tests is disabled.		
TEST			
DEXPT	Field	Value	Comments
EWASC	TEST	0	Failure prediction processing and warning processing shall be disabled.
	DEXCPT	1	
	EWASC	0	
EBF	Shall be set to 0b indicating device specific background functions are disabled.		
PERF	Shall be set to 1b indicating processing of informational exception conditions that cause delays are disabled.		
MRIE	Shall be set to 0h indicating device server shall not report informational exception conditions.		
INTERVAL TIMER	Shall be set to zero indicating reporting informational exception conditions are vendor specific.		
REPORT COUNT	Shall be set to zero indicating no limit on number of times a device server may report informational exception conditions.		

6.4 READ CAPACITY Parameter Data

READ CAPACITY Parameter Data is returned to application client in response to a READ CAPACITY(10) or READ CAPACITY(16) command (refer to 5.4).

Table 6-27: Translation of READ CAPACITY Parameter Data

READ CAPACITY Parameter Data Field	Notes and References											
RETURNED LOGICAL BLOCK ADDRESS	Shall be set to:											
	READ CAPACITY (10)	Shall be set to Namespace Size NSZE field of the Identify Namespace Data Structure, or FFFF_FFFFh if NSZE exceeds the maximum value this field is able to contain. Note that if this field is set to FFFF_FFFFh then SBC-3 specifies that the application client issue a READ CAPACITY (16).										
	READ CAPACITY(16)	Shall be set to Namespace Size NSZE field of the Identify Namespace Data Structure, or FFFF_FFFF_FFFF_FFFFh if NSZE exceeds the maximum value this field is able to contain.										
LOGICAL BLOCK LENGTH	Shall be set to the LBA Data Size (LBADS) field of the LBA Format Data Structure currently indicated by the Formatted LBA Size (FLBAS) field within the Identify Namespace Data Structure.											
PROT_EN*	Shall be set to 0b if protection information is not enabled as indicated in the End-to-end Protection Type Settings (DPS) field within Identify Namespace Data Structure. Otherwise shall be set to 1b.											
P_TYPE*	Shall be set using End to End Data Protection Type Settings field of Identify Namespace Data Structure: <table><tr><th>DPS Code</th><th>P_TYPE Code</th></tr><tr><td>000b</td><td>Unspecified</td></tr><tr><td>001b</td><td>000b</td></tr><tr><td>010b</td><td>001b</td></tr><tr><td>011b</td><td>010b</td></tr></table>		DPS Code	P_TYPE Code	000b	Unspecified	001b	000b	010b	001b	011b	010b
DPS Code	P_TYPE Code											
000b	Unspecified											
001b	000b											
010b	001b											
011b	010b											
LOGICAL BLOCKS PER PHYSICAL EXPONENT*	Shall be set to 0 indicating one or more physical blocks per logical block.											
P_I_EXPONENT*	Unspecified. Indicates number of protection information intervals within each logical block.											
TPE*	Unspecified. Indicates thin provisioning is implemented.											
TPRZ*	Unspecified. Indicates how unmapped LBA bits should be set.											
LOWEST ALIGNED LOGICAL BLOCK ADDRESS*	Shall be set to zero indicating LBA 0 is the first logical blocks located at the beginning of a physical block.											
*READ CAPACITY (16) parameter data only												

6.5 Sense Data

Sense data is requested by the application client with a REQUEST SENSE command. Since autosense is supported, sense information is returned along with the completion of a command to the application client.

6.5.1 Fixed Format Sense Data

Fixed Format Sense Data Field	Notes and References
RESPONSE CODE	Shall be set to 70h indicating current errors are returned.
VALID	Shall be set to 0b indicating INFORMATION field is not defined.
SENSE KEY	Shall be set to NO SENSE.
ILI	Shall be set to 0b, field is undefined.
EOM	Shall be set to 0b, field is undefined.
FILEMARK	Shall be set to 0b, field is undefined.
INFORMATION	Shall be set to 0b, field is undefined.
ADDITIONAL SENSE LENGTH	Shall indicate the number of bytes of additional sense data.
COMMAND-SPECIFIC INFORMATION	Shall contain information depending on command where the exception condition occurred.
ADDITIONAL SENSE CODE	Shall be set to NO ADDITIONAL SENSE INFORMATION if device is in power state 00h, otherwise shall be set to LOW POWER CONDITION ON.
ADDITIONAL SENSE CODE QUALIFIER	Shall be set to zero indicating detailed information is not available.
FIELD REPLACEABLE UNIT CODE	Shall be set to zero indicating no additional sense code.
SENSE KEY SPECIFIC	Shall be set to zero, field is undefined.
SKSV	Shall be set to zero indicating SENSE KEY SPECIFIC field is undefined.

6.5.2 Descriptor Format Sense Data

Table 6-28: Descriptor Format Sense Data Field Translations

Fixed Format Sense Data Field	Notes and References
RESPONSE CODE	Shall be set to 72h indicating current errors are returned.
SENSE KEY	Shall be set to NO SENSE.
ADDITIONAL SENSE CODE	Shall be set to NO ADDITIONAL SENSE INFORMATION if device is in power state 00h, otherwise shall be set to LOW POWER CONDITION ON.
ADDITIONAL SENSE CODE QUALIFIER	Shall be set to zero indicating detailed information is not available.
ADDITIONAL SENSE CODE LENGTH	Shall be set to zero indicating no additional sense code.

6.6 REPORT LUNS Parameter Data

REPORT LUNS Parameter data is requested by the application client with a REPORT LUNS command (refer to 4.5). Support for REPORT LUNS requires the Identify command.

Table 6-29: REPORT LUNS Parameter Data Field Translations

REPORT LUNS Parameter Data Field	Notes and References
LUN LIST LENGTH	Shall indicate the length of the LUN LIST as specified in SPC-4. The LUN LIST LENGTH depends on the translation of the LUN LIST field.
LUN LIST	The list shall contain logical unit numbers corresponding to namespaces present on the device with a Namespace Capacity (NCAP) field of the Identify Namespace Structure set to greater than 0h. Logical unit numbers shall begin with 0 and have a maximum value of NN-1, where NN is the Number of Namespaces field within Identify Controller Data Structure.

6.7 PERSISTENT RESERVE IN/OUT Parameter Data

Many translations require the SNTL send NVM Reservation Report command, which requires a buffer allocated to transfer the report, and report that buffer's size in DWORDs. The minimum size in DWORDs the host should have allocated for the Reservation Status buffer is 12, with 6 additional for every registered controller in the NVMe subsystem. The NVMe Reservation Status data structure's REGCTL can be used to calculate how many DWORDS should have been allocated if the buffer in the original request was not large enough.

6.7.1 PERSISTENT RESERVE IN Parameter Data for READ KEYS

The SNTL shall support this SCSI command by issuing NVM Reservation Report (0Eh) command with the NSID set for the namespace being queried with command DWORD 10 (NUMD) set to the 0's based size of the Reservation Status structure, and then translating from the following table.

Table 6-30: READ KEYS

READ KEYS Field	Notes and References
PRGENERATION	Shall be set to Reservation Status Data Structure GEN field, returned from Reservation Report command.
ADDITIONAL LENGTH	Shall be set to the number of additional bytes that follow, which is the number of registered keys (Reservation Status Data Structure REGCTL field) times the size of each key (8 bytes).
RESERVATION KEY [first]	Shall be set to Reservation Status Data Structure RKEY field from the first controller returned from Reservation Report command.
...	

RESERVATION KEY [last]	Shall be set to Reservation Status Data Structure RKEY field from the last controller returned from Reservation Report command.
------------------------	---

6.7.2 PERSISTENT RESERVE IN Parameter Data for READ RESERVATIONS

The SNTL shall support this SCSI command by issuing NVM Reservation Report (0Eh) command with NSID set to the namespace being queried and command DWORD 10 (NUMD) set to the 0's based size of the buffer used for the Reservation Status structure, and then translating from the following table.

Table 6-31: READ RESERVATIONS

READ RESERVATIONS Field	Notes and References
PRGENERATION	Shall be set to Reservation Status Data Structure GEN field, returned from Reservation Report command.
ADDITIONAL LENGTH	Shall be set to Reservation Status Data Structure 10h, or 00h if no reservation is held (SPC-4).
RESERVATION KEY	Shall be set to Reservation Status Data Structure RKEY field from the controller holding the reservation from Reservation Report command.
SCOPE	Shall be set to 0h (SPC-4).
TYPE	Shall be translated from Reservation Status RTYPE field. See Table 4-14.

6.7.3 PERSISTENT RESERVE IN Parameter Data for READ CAPABILITIES

The SNTL shall support this SCSI command by issuing NVMe Identify (06h) command with CNS set to 0 for Identify Namespace and NSID set to the namespace being queried. The SNTL shall also issue Get Features (09h) with FID set to Reservation Persistence 83h. Once both commands are completed, the translation is done as defined in the following table.

Table 6-32: READ CAPABILITIES

READ CAPABILITIES Field	Notes and References
LENGTH	Set to 0008h (SPC-4)
CRH	Compatible Reservation Handling shall set to 0b as SPC-2 RESERVE commands won't be handled at all and shall be terminated with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID OP CODE additional sense.
SIP_C	Specify Initiation Ports Capability shall be set to 0b as there is not NVMe parameter to specify a specific host.

ATP_C	All Target Port Capability shall be set to 1b as the NVMe subsystem applies registration to all controllers with matching the host identifier (HOSTID).														
PTPL_C	Persistent through Power Loss Capability shall be set to bit 0 from Identify Namespace's RESCAP field.														
PTL_A	Persistent Through Power Loss Activated shall be set to Get Feature "Reservation Persistence" (Feature Identifier 83h) PTPL value from the Command Completion's command specific status field.														
ALLOWED COMMANDS	Shall be set to 000b, making no guarantees on allowed commands.														
TMV	Type Mask Valid shall be set 1 to indicate the reservation type mask is valid.														
PERSISTENT RESERVATION TYPE MASK	<p>From Identify Namespace, translate the returned value in RESCAP field as defined in the following table:</p> <table> <tr> <th>Type</th><th>Notes</th></tr> <tr> <td>WR_EX</td><td>Shall be set if RESCAP bit 1 is set.</td></tr> <tr> <td>EX_AC</td><td>Shall be set is RESCAP bit 2 is set.</td></tr> <tr> <td>WR_EX_RO</td><td>Shall be set if RESCAP bit 3 is set.</td></tr> <tr> <td>EX_AC_RO</td><td>Shall be set if RESCAP bit 4 is set.</td></tr> <tr> <td>WR_EX_AR</td><td>Shall be set if RESCAP bit 5 is set.</td></tr> <tr> <td>EX_AC_AR</td><td>Shall be set if RESCAP bit 6 is set.</td></tr> </table>	Type	Notes	WR_EX	Shall be set if RESCAP bit 1 is set.	EX_AC	Shall be set is RESCAP bit 2 is set.	WR_EX_RO	Shall be set if RESCAP bit 3 is set.	EX_AC_RO	Shall be set if RESCAP bit 4 is set.	WR_EX_AR	Shall be set if RESCAP bit 5 is set.	EX_AC_AR	Shall be set if RESCAP bit 6 is set.
Type	Notes														
WR_EX	Shall be set if RESCAP bit 1 is set.														
EX_AC	Shall be set is RESCAP bit 2 is set.														
WR_EX_RO	Shall be set if RESCAP bit 3 is set.														
EX_AC_RO	Shall be set if RESCAP bit 4 is set.														
WR_EX_AR	Shall be set if RESCAP bit 5 is set.														
EX_AC_AR	Shall be set if RESCAP bit 6 is set.														

6.7.4 PERSISTENT RESERVE IN Parameter Data for READ FULL STATUS

The SNTL shall support this SCSI command by issuing NVM Reservation Report (09h) command with NSID set to the namespace being queried and command DWORD 10 (NUMD) set to the 0's based size of the reservation Status structure, and then translating in the following table.

Table 6-33: READ FULL STATUS

READ FULL STATUS Field	Notes and References
PRGENERATION	Shall be set to Reservation Status Data Structure GEN field, returned from Reservation Report command.

ADDITIONAL LENGTH	Shall be set to the number of bytes that follow, which is the number of registered keys (Reservation Status Data Structure REGCTL field) times the size of each descriptor (32 bytes, defined in Table 6-34).
FULL STATUS DESCRIPTOR [first]	Shall be set from Reservation Status Data Structure field from the first controller returned from Reservation Report command (Table 6-34).
...	
FULL STATUS DESCRIPTOR [last]	Shall be set from Reservation Status Data Structure field from the last controller returned from Reservation Report command (Table 6-34).

Table 6-34: FULL STATUS DESCRIPTOR

FULL STATUS DESCRIPTOR Field	Notes and References
RESERVATION KEY	Shall be set to this entry's RKEY value from NVME Reservation Report Registered Controller Data Structure.
R HOLDER	Shall be set to 1b if this entry's NVME Reservation Report Registered Controller Data Structure RCSTS field bit 1 is set, and 0b otherwise.
ALL_TG_PT	Shall be set to 1b.
SCOPE	Shall be set to 00h.
TYPE	Set from this entry's RTYPE value. See Table 4-14.
RELATIVE TARGET PORT ID	Shall be set to this entry's CNTLID from NVME Reservation Report Registered Controller Data Structure.
ADDITIONAL DESCRIPTOR LENGTH	Shall be set to 08h to account for the Transport ID.
TRANSPORT ID	Shall be set to this entry's HOSTID from NVMe Reservation Report Registered Controller Data Structure.

6.7.5 PERSISTENT RESERVE OUT Parameter List

The SNLT shall support this SCSI command by issuing NVM Reservation Acquire (11h), Register (0Dh), or Release (15h) command depending on the service action. The specific command to issue depends on the

SCSI command's Service Action from CDB (see 4.13.2 for details), and translation to each action and parameter list fields are defined in the following table.

Table 6-35: PERSISTENT RESERVE OUT

PARAMETER LIST	Notes and References	
RESERVATION KEY	Service Action	Notes
	REGISTER AND MOVE	<p>The SNTL shall set the Reservation Register Data Structure CRKEY to the parameter list RESERVATION KEY.</p> <p>Reservation Register command's DWORD 10 RREGA field shall be set to 010b (Replace).</p>
	RESERVE/PREEMPT/PREEMPT AND ABORT	The SNTL shall set the Reservation Acquire Data Structure CRKEY to the parameter list RESERVATION KEY.
	RELEASE/CLEAR	The SNTL shall set the Reservation Release Data Structure CRKEY to the parameter list RESERVATION KEY.
	REGISTER/REGISTER AND IGNORE EXISTING KEY	<p>If SERVICE ACTION RESERVATION KEY is zero, the SNTL shall set the Reservation Register Data Structure CRKEY to parameter list RESERVATION KEY, and shall set Reservation Register command's DWORD 10 RREGA field 001b (Unregister).</p> <p>If SERVICE ACTION RESERVATION KEY is non-zero, this parameter field shall be ignored, the SNTL shall treat the Reservation Register Data Structure CRKEY as reserved, and shall set Reservation Register command's DWORD 10 RREGA field to 000b (Register).</p>
SERVICE ACTION RESERVATION KEY	If this field is 0, it is ignored. If the field is non-zero, the translation is in the following table:	

	Service Action	Notes
	REGISTER/REGISTER AND IGNORE EXISTING KEY/REGISTER AND MOVE	The SNTL shall set the Reservation Register Data Structure NRKEY to the parameter list SERVICE ACTION RESERVATION KEY.
	PREEMPT/ PREEMPT AND ABORT	The SNTL shall set the Reservation Acquire Data Structure PRKEY to the parameter list SERVICE ACTION RESERVATION KEY.
	RESERVE/RELEASE/CLEAR	Unused
SPEC_I_PT	The SNTL shall only support translating this command if SPEC_I_PT parameter is set to 0 as reservations are applied on an NVMe Subsystem and not per I_T_Nexus; if set to 1, the SNTL shall terminate the command with CHECK CONDITION status, ILLEGAL REQUEST sense key, and INVALID FIELD IN PARAMETER LIST additional sense.	
ALL_TG_PT	Ignored. NVMe Subsystem applies registration action to all controllers with the associated HOSTID.	
APTPL	Service Action	Notes
	REGISTER/REGISTER AND IGNORE EXISTING KEY	<p>If APTPL is set to 0b, Reservation Register command DWORD 10 CPTPL shall be set to 10b to request reservations and registrants are cleared on power on.</p> <p>If APTPL is set to 1b, Reservation Register command DWORD 10 CPTPL shall be set to 11b to request reservations and registrants are persistent across power loss if the device supports this (refer to table 6-28, field PTPL_C). If the device does not support persisting reservations through power loss, the command shall be terminated with CHECK CONDITION status,</p>

		<p>ILLEGAL REQUEST sense key, and INVALID FIELD IN PARAMETER LIST additional sense.</p> <p>If APTPL is set to 1b and device supports persistent reservations through power lost, the SNTL shall issue a Set Features Admin command for Reservation Persistence Feature ID 83h, with command DWORD 11 bit 0 (PTPL) set to 1 if it is not already set (refer to table 6-31, field PTPL_A) prior to issuing the Reservation Register NVM command.</p>
	<p>PREEMPT/ PREEMPT AND ABORT/ RESERVE/RELEASE/CLEAR/ REGISTER AND MOVE</p>	Unused.
ADDITIONAL PARAMETER DATA	Unused	

7 Status Mappings

Command statuses shall require translation from a NVM Express Status Code to a SCSI Status Code, Sense Key, and an Additional Sense Code where applicable.

7.1 Generic Command Status

Table 7-1: Generic Command Status Mappings

<i>NVM Express</i>	<i>SCSI</i>		
Status Code	Status Code	Sense Key	Additional Sense Code
Success Completion	GOOD	NO SENSE	
Invalid Command Opcode	CHECK CONDITION	ILLEGAL REQUEST	INVALID COMMAND OPERATION CODE
Invalid Field in Command	CHECK CONDITION	ILLEGAL REQUEST	INVALID FIELD IN CDB
Data Transfer Error	CHECK CONDITION	MEDIUM ERROR	
Commands Aborted due to Power Loss Notification	TASK ABORTED	ABORTED COMMAND	WARNING – POWER LOSS EXPECTED
Internal Device Error	CHECK CONDITION	HARDWARE ERROR	INTERNAL TARGET FAILURE
Command Abort Requested	TASK ABORTED	ABORTED COMMAND	
Command Aborted due to SQ Deletion	TASK ABORTED	ABORTED COMMAND	
Command Aborted due to Failed Fused Command	TASK ABORTED	ABORTED COMMAND	
Command Aborted due to Missing Fused Command	TASK ABORTED	ABORTED COMMAND	
Invalid Namespace or Format	CHECK CONDITION	ILLEGAL REQUEST	ACCESS DENIED – INVALID LU IDENTIFIER
LBA Out of Range	CHECK CONDITION	ILLEGAL REQUEST	LOGICAL BLOCK ADDRESS OUT OF RANGE
Capacity Exceeded	CHECK CONDITION	MEDIUM ERROR	
Namespace Not Ready (Do Not Retry bit value 1)	CHECK CONDITION	NOT READY	LOGICAL UNIT NOT READY, CAUSE NOT REPORTABLE
Namespace Not Ready (Do Not Retry bit value 0)	CHECK CONDITION	NOT READY	LOGICAL UNIT NOT READY, BECOMING READY
Reservation Conflict	RESERVATION CONFLICT	N/A	N/A

7.2 Command Specific Status

Table 7-2: Command Specific Status Mappings

<i>NVM Express</i>	<i>SCSI</i>		
Status Code	Status Code	Sense Key	Additional Sense Code
Completion Queue Invalid	CHECK CONDITION	ILLEGAL REQUEST	
Invalid Queue Identifier	Unspecified	Unspecified	Unspecified
Maximum Queue Size Exceeded	Unspecified	Unspecified	Unspecified
Abort Command Limit Exceeded	CHECK CONDITION	ILLEGAL REQUEST	
Asynchronous Event Request Limit Exceeded	Unspecified	Unspecified	Unspecified
Invalid Firmware Slot	Unspecified	Unspecified	Unspecified
Invalid Firmware Image	Unspecified	Unspecified	Unspecified
Invalid Interrupt Vector	Unspecified	Unspecified	Unspecified
Invalid Log Page	Unspecified	Unspecified	Unspecified
Invalid Format	CHECK CONDITION	ILLEGAL REQUEST	FORMAT COMMAND FAILED
Conflicting Attributes	CHECK CONDITION	ILLEGAL REQUEST	INVALID FIELD IN CDB
Attempted Write to Read Only Range	CHECK CONDITION	DATA PROTECT	WRITE PROTECTED

7.3 Media Errors

Table 7-3: Media Errors Mappings

<i>NVM Express</i>	<i>SCSI</i>		
Status Code	Status Code	Sense Key	Additional Sense Code
Write Fault	CHECK CONDITION	MEDIUM ERROR	PERIPHERAL DEVICE WRITE FAULT
Unrecovered Read Error	CHECK CONDITION	MEDIUM ERROR	UNRECOVERED READ ERROR
End-to-end Guard Check Error	CHECK CONDITION	MEDIUM ERROR	LOGICAL BLOCK GUARD CHECK FAILED
End-to-end Application Tag Check Error	CHECK CONDITION	MEDIUM ERROR	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
End-to-end Reference Tag Check Error	CHECK CONDITION	MEDIUM ERROR	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
Compare Failure	CHECK CONDITION	MISCOMPARE	MISCOMPARE DURING VERIFY OPERATION
Access Denied	CHECK CONDITION	DATA PROTECT	ACCESS DENIED – NO ACCESS RIGHTS

8 Task Management Functions

Task Management Functions are requests by application clients affecting the processing of a SCSI commands. Support for Task Management Functions requires maintaining a mapping between the SCSI command issued by the application client and one or more NVM Express commands resulting from translation into NVM Express.

Table -8-1: Translation of Task Management Functions

Function	Notes
ABORT TASK	May be supported by issuing an Abort command.
ABORT TASK SET	May be supported by issuing an Abort command on all outstanding commands.
CLEAR ACA	Unspecified
CLEAR TASK SET	May be supported by issuing an Abort command on all outstanding commands.
I_T NEXUS RESET	Shall be supported by returning FUNCTION SUCCEEDED if there are outstanding commands in the submission queue, otherwise by returning FUNCTION COMPLETE.
LOGICAL UNIT RESET	Shall be supported by writing a 0 to Enable (EN) field of Controller Configuration register.
QUERY TASK	May be supported. If specified command has one or more NVM Express commands queued resulting from translation than FUNCTION SUCCEEDED shall be returned. Otherwise, FUNCTION COMPLETE shall be returned.
QUERY TASK SET	Unspecified
QUERY ASYNCHRONOUS EVENT	Unspecified
